Project no.:	101095738		
Project full title:	6G SHort range extreme communication IN Entities		
Project Acronym:	6G-SHINE		
Project start date:	01/03/2023	Duration	30 months

D5.3 – SECOND IMPLEMENTATION AND EVALUATION OF POCS, FINAL RESULTS

Due date	30/08/2025	Delivery date	29/08/2025
Work package	WP5		
Responsible Author(s)	Spilios Giannoulis (IMEC)		
Contributor(s)	Pedro Maia de Sant Ana (BOSCH), Wei Liu (IMEC), Thijs Havinga (IMEC), Anastasios Grammenos (COGN), Christos Tsakos (COGN), Fotis Foukalas (COGN), Oliver Haala (FHG), Jürgen Hupp (FHG), Spilios Giannoulis (IMEC), Irfan Jabandzic (IMEC), Filipe Conceicao, Sebastian Robitzsch, Muhammad Awais Jadoon, Ayesha Ijaz (IDE), Usman Virk (Keysight)		
Version	V1.0		
Reviewer(s)	Baldomero Coll Perales (UMH), Gilberto Berardinelli (AAU)		
Dissemination level	Public		

Co-funded by the European Union

Horizon Europe Grant Agreement No. 101095738. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or SNS JU. Neither the European Union nor the granting authority can be held responsible for them.

VERSION AND AMENDMENT HISTORY

Version	Date (DD/MM/YYYY)	Created/Amended by	Changes
0.1	12/05/2025	Spilios Giannoulis	TOC finalization
0.2	30/05/2025	Spilios Giannoulis	Initial contributions from partners
0.3	09/06/2025	Spilios Giannoulis	Editor revision and feedback
0.4	27/06/2025	Spilios Giannoulis	Refined contributions
0.5	02/07/2025	Spilios Giannoulis	Editor review
0.6	04/07/2025	Spilios Giannoulis	Final version ready for internal review
0.7	17/08/2025	Spilios Giannoulis	Final editors' version
0.9	28/08/2025	Berit H. Christensen	Proofreading and layout check completed
1.0	29/08/2025	Spilios Giannoulis	Final version

TABLE OF CONTENTS

FIGURES LIST

TABLES LIST

ABBREVIATIONS

Acronym	Description
AAL	Antenna array-like
AD	Analog devices
AI	Artificial intelligence
AO	Alternate optimization
AP	Access point
BACKCOM	Backscatter communication
BCD	Block coordinate descent
BS	Base station
CCA	Clear Channel Assessment
COT	Channel Occupancy Time
CR	Constraint relaxation
CRUA	Clear Resource Unit Assessment
CSI	Channel state information
CSMA/CA	Carrier Sense Multiple Access / Collision Avoidance
CTI	Cross Technology Interference
CW	Continuous wave
DC	Dual connectivity
DMA	Dynamic metasurface antenna
DUT	Device-under-test
EM	Electromagnetic
ETSI	European Telecommunications Standards Institute
EW	Electronic warfare
FDMA	Frequency division multiple access
FPGA	Field programmable gate array
FSM	Finite state machine
HC	High-capacity device
IEEE	Institute of Electrical and Electronics Engineers
IO	Indoor office
IOT	Internet-of-things
ISM	Industrial, Scientific and Medical
KKT	Karush-Kuhn-Tucker
LA	Latency Aware
LOS	Line-of-sight
LR-WPAN	Low rate wireless personal area network
MAC	Medium Access control
MIMO	Multiple-input multiple-output
MISO	Multiple input single output
ML	Machine learning
MSE	Mean square error
MU	Multi user

NLOS	Non-line-of-sight
NOMA	Non orthogonal multiple access
N-RIS	Non-reconfigurable intelligent surface
OFDM	Orthogonal frequency division multiplexing
OFDMA	Orthogonal frequency division multiple access
PDP	Power-delay profile
PER	Packet Error Rate
PoC	Proof of Concept
RIS	Reconfigurable Intelligent Surface
RRM	Radio Resource Management
RU	Resource unit
SDR	Software defined radio
SNR	Signal to noise ratio
SU	Single user
SW	Software
TC	Technology Component
TCF	Temporal correlation function
VNA	Vector network analyser

EXECUTIVE SUMMARY

Deliverable 5.3 is the final deliverable of WP5, where a proof of concept (PoC) demonstration of selected 6G-SHINE technologies is developed. D5.3 takes input from Tasks 5.1 and 5.2 “Detailed presentation of the technology components to be demonstrated” and “Initial implementation of PoC components” of this project as well as from the TC WPs (WP2, WP3 and WP4) where enhanced designs of the TCs took place during the project timeline. In this task, each PoC is enhanced according to a commonly agreed methodology, considering the results achieved in T5.2 as well as lessons learned from the further analytical or simulation evaluation of the relevant TCs, helping to define the list of added or enhanced features of the TCs implemented into PoCs at this final stage. Relevant key performance indicators (KPIs) are used to evaluate the performance for each PoC and where possible compared to the previous PoC version and/or the theoretical or simulation results available from the relevant TC’s WP.

The final list of PoCs presented in this deliverable are listed below, along with their connection to the respective technology components (TCs) defined in the 6G-SHINE project proposal:

- Low latency channel emulator (addressing TC2)
- Latency aware MAC access (addressing TC11)
- Jamming resilient PHY (addressing TC6)
- Intra-subnetwork macro-diversity (addressing TC8)
- Centralized RRM (addressing TC13)
- Centralized/distributed interference management (address both TC12 and TC13)
- Jamming detection using anomaly detection (addressing TC15)

For each PoC, we present the final implementation state, architecture, hardware and software building blocks, applicable KPIs and demonstration scenario. In addition, a full result presentation and evaluation segment is presented for each PoC where the achieved results are also compared with the defined baseline and the results achieved in T5.2 and in the theoretical WPs of each PoC and its relevant TC in WP2,3,4.

1 INTRODUCTION

The 6G-SHINE project's Work Package 5 (WP5) is dedicated to the realization and validation of proof-of-concept (PoC) demonstrations for the in-X wireless subnetworks technologies under development. Building on the theoretical and simulation studies carried out in WP2, WP3 and WP4, WP5's main objective is to identify the most promising technical solutions for targeted use cases, to implement and evaluate an initial set of PoC components, and then to refine those implementations toward more challenging and realistic scenarios. This iterative approach ensures that the consortium can both verify key performance indicators (KPIs) in laboratory conditions and provide actionable feedback to the design-oriented work packages (WP2–WP4) to optimize their solutions for real-world deployment.

6G-SHINE is an exploratory project focused on developing foundational technology components for in-X subnetworks, spanning from the physical layer to architectural innovations. A total of 16 technology components (TCs) were identified as relevant to this goal. Among them, 7 were highlighted in the grant agreement as candidates for proof-of-concept (PoC), based on their expected maturity level - specifically, a Technology Readiness Level (TRL) of 4, meaning validation in a laboratory environment.

As outlined in previous deliverables, the selected PoC candidates are:

1. **Low Latency Channel Emulator (Keysight)** – validating **TC2**: Channel models for in-X scenarios
2. **Latency-Aware MAC (IMEC)** – validating **TC11**: Latency-aware access in unlicensed spectrum
3. **Jamming-Resilient PHY (IMEC)** – validating **TC6**: Jamming-aware native PHY design
4. **Intra-Subnetwork Macro-Diversity (FHG)** – validating **TC8**: Intra-subnetwork macro-diversity
5. **Wireless In-Vehicle E/E Architecture (Bosch)** – validating **TC12**: Centralized radio resource management
6. **Centralized/Distributed Interference Management (COGN)** – validating **TC12** and **TC13**: Centralized and distributed/hybrid radio resource management
7. **Jamming Detection via Anomaly Detection Methods (IDE)** – validating **TC15**: Hybrid management of traffic, spectrum, and computational resources

The remaining technology components were investigated up to TRL 2, through theoretical analyses and simulation-based studies reported in Work Packages 2, 3, and 4.

Deliverable D5.2 ("First implementation and evaluation of PoCs") [1] captured the outcomes of WP5's initial implementation phase. In that phase, partners deployed each PoC's hardware and software in their respective testbeds, defined experimental setups and orchestration procedures, and measured performance against baseline employing related KPIs per PoC (e.g., packet error rate, latency, spectral efficiency, energy consumption, etc.). The results were critically benchmarked against WP3/WP4 simulations, revealing both the practical strengths of the proposed designs and the areas where hardware or signalling constraints required further adaptation. These findings laid the groundwork for more advanced PoC iterations.

This deliverable (D5.3, "Second implementation and evaluation of PoCs, final results") presents the refined PoC implementations informed by D5.2's lessons learned. The consortium integrated evolved design features -such as denser deployment handling and non-ideal signalling adaptations -into the existing testbed platforms. Performance has been reevaluated against the baseline and, where applicable, enhanced KPI sets under more challenging operational scenarios that more closely mirror end-user environments. The presented final results are used to validate the in-X subnetwork enablers beyond simulation, highlight any remaining limitations, and demonstrate each technology's readiness for the next phase of the SNS-JU work program.

More specifically, in Chapter 2 we focus on presenting the final results coming out of each PoC following a unified presentation approach, when possible, between the different PoCs. Changes in architecture, in functionality and demonstrator setup or scenarios are clearly described and explained in detail while updated results are presented against the defined baseline or the previous achieved results coming from D5.2 [1]. In Chapter 3 the connection between WP5 PoCs and the relevant research conducted in WP2, WP3 and WP4 is presented, in order to become

clear where the PoCs are coming from, what is their theoretical background and also how they are used to verify theoretical or simulation results from the mentioned WPs. Finally, in Chapter 4, conclusions are drawn based on presented results and future steps are described where possible.

2 PROOF OF CONCEPT PRESENTATION

In this chapter, we present the second iteration of implemented PoCs in detail, focusing on the enhancements presented in line with the relevant work in WP2, WP3 and WP4. Each PoC presentation consists of a general description, followed by a detailed description of its enhanced features based on what was already presented in D5.2 in terms of functionality and/or architecture design as well as demonstrator scenario. Finally, the evaluation results are presented and compared to their defined baseline as well to the related theoretical or simulation results presented in the relevant project WP where the initial TC was presented when possible. The relevance of each PoC for chosen use-cases is presented per PoC as well as its quantitative impact against the defined KPIs of the chosen use case. This approach will make clear why but also how each PoC can really act as an enabler for the realization of the use cases envisioned in WP2.

2.1 LOW-LATENCY CHANNEL EMULATOR

2.1.1 GENERAL DESCRIPTION

Wireless communication systems face significant challenges due to propagation phenomena such as attenuation, shadowing, fast fading, multipath delays, Doppler effects, and interference. These effects can severely degrade reliability, throughput, and timing accuracy of transmissions. Traditional field-testing approaches, though valuable, are often labour-intensive, expensive, time-consuming, and lack repeatability due to constantly changing environmental conditions. To overcome these limitations, radio channel emulators have become indispensable for the development and validation of advanced wireless technologies. These emulators provide a controlled laboratory environment that enables deterministic, rigorous, and repeatable testing of radio systems under real-world conditions.

The 6G-SHINE project targets future communication environments that demand ultra-reliable low-latency communication (URLLC), particularly in short-range "in-X subnetworks" such as those found in industrial automation, vehicular systems, and immersive consumer experiences. These use cases require extremely short propagation delays, often below 1 µs, to ensure synchronized operation of time-critical applications. Achieving such latency constraints in a test environment necessitates significant advancements in radio channel emulator architecture and design. To understand the importance of low-latency design, consider that a 1 µs delay in the emulator represents a physical distance of 300 meters in free-space propagation. With emerging 6G systems expected to operate at much shorter ranges and higher frequencies, large delays become unacceptable, often causing synchronization failures in time-critical applications. Thus, reducing minimum emulator latency is imperative for meeting the requirements of evolving 6G use cases.

This section outlines the technical concept, implementation advancements, architecture updates, and performance results of the improved PoC and continues the work presented in D5.1 and D5.2 [1] by documenting and demonstrating the enhancements of Keysight's low-latency radio channel emulator platform, PROPSIM. The original proof-of-concept (PoC) demonstrated in D5.2 [1] focused on software and FPGA-level optimizations to reduce insertion delay (latency) to sub-microsecond levels (~0.9 µs). D5.3 now extends this effort by incorporating hardware architectural modifications and next-generation ADC/DAC designs that target latency reductions down to 0.5 µs and beyond.

The enhanced emulator is capable of replicating real-world wireless conditions, including multipath propagation, delay spread, noise, and Doppler effects. It supports a wide array of standards and configurations, including LTE, NR, Wi-Fi, 3GPP, CTIA, MIMO, beamforming, and aerospace communications, respectively. Its vendor-agnostic, file-based channel emulation approach supports customizable channel models with high configurability. Built-in tools and a user-friendly GUI make the emulator suitable for both basic lab setups and advanced R&D investigations.

This work builds directly on the theoretical and modelling efforts of WP2 D2.3 [1], which characterized and developed the short-range channel models and corresponding KPIs for in-X subnetworks, mainly for the consumer and industrial use cases. In WP5, the implementation of these models in an emulator context has enabled practical

validation of latency-critical scenarios under repeatable test conditions. The work in D5.3 aligns with the broader 6G-SHINE objectives of enabling time-synchronized and performance-optimized communication systems.

2.1.2 ENHANCED/NEW FUNCTIONALITY VS PREVIOUS VERSION

The deliverable D5.2 [1] introduced the first Proof-of-Concept (PoC) implementation of a low-latency radio channel emulator, based on Keysight's PROPSIM platform, and targeted for 6G in-X subnetwork applications. That implementation focused on software-level and FPGA optimizations to reduce end-to-end insertion delay from the baseline 2.2 µs to ~1.0 µs, without modifying the underlying hardware.

The improvements in D5.2 [1] primarily involved:

1. Optimization of component settings, including buffering, clock domain alignment, and signal timing.
2. FPGA-level enhancements to critical DSP blocks, such as the DDC, DUC, equalizer (EQU), and fading module.

A range of KPIs, derived from the industrial use case channel models developed in WP2 Task 2.2, D2.3 [1], were successfully validated:

1. Insertion Delay: Reduced to ~0.9 µs using S21-based delay measurements with a Keysight PNA-X
2. Power Delay Profile (PDP): Validated by reconstructing impulse responses from frequency-domain measurements
3. Doppler Spectrum and Temporal Correlation Function (TCF): Measured using Keysight UXA and analysed offline in MATLAB

These results confirmed the emulator's ability to accurately replicate real-world channel dynamics under latency-critical conditions, establishing a strong baseline for the hardware-level innovations introduced in D5.3.

2.1.2.1 FUNCTIONAL ENHANCEMENTS BEYOND D5.2

The current work, presented in D5.3, builds upon the WP2-defined latency-critical use cases, such as industrial automation, vehicular safety systems, and immersive XR environments, which require deterministic timing behaviour and sub-microsecond propagation delay to maintain synchronization and operational reliability.

In alignment with Technology Component 2 (TC2), i.e. Channel models for in-X scenarios, this updated PoC targets the following key functional advancements:

1. **TC2-A1:** Latency reduction to ≤ 0.5 µs, meeting the targets set in 6G-SHINE
2. **TC2-A2:** Integration of next-generation ADC/DAC components with embedded DSP
3. **TC2-A3:** Complete re-architecture of the baseband DSP processing chain
4. **TC2-A4: Low-latency fading block:** Deployment of a new low-latency fading block
5. **TC2.A5:** Introduction of a latency-optimized routing and combining subsystem

TABLE 1. SUMMARY OF ENHANCEMENTS COMPARED TO THE PREVIOUS VERSION

Feature	D5.2 Implementation	D5.3 Improvements
ADC/DAC Subsystem	External FPGA-based DDC/DUC/EQ	Embedded DDC/DUC/EQ within ADC/DAC; higher sampling clocks; reduced conversion latency
Fading Block	Legacy implementation in standard FPGA fabric	Re-implemented using hardened DSP slices; latency reduced to < 50 ns
DSP Chain Architecture	Multistage logic with serialized interconnects	Re-architected with fewer buses and synchronized clock domains
Interconnect Interface	Traditional bus-based coupling	JESD204-based high-speed serial interfaces between ADC/DSP domains

These enhancements advance the emulator into a new performance regime while preserving functional integrity across critical channel model metrics, including PDP, Doppler spectrum, and temporal correlation. The architectural redesign ensures the emulator can support realistic, real-time 6G in-X subnetwork testing scenarios with deterministic low-latency performance. In Table 1, the specific advancements vs the previous version of the PoC is presented and explained in detail.

2.1.3 HIGH LEVEL ARCHITECTURE CHANGES/UPDATES

The updated Proof-of-Concept (PoC) demonstrator in D5.3 introduces significant architectural and hardware modifications aimed at achieving an end-to-end latency target of less than equal to 0.5 µs. These enhancements build upon the software and FPGA-based optimizations developed in D5.2, extending them with hardware innovations including new high-speed ADC/DAC modules, optimized DSP chains, and re-engineered data routing.

2.1.3.1 REVISED SYSTEM ARCHITECTURE COMPONENTS

The PoC demonstrator now comprises the following key subsystems, as it can be seen in Figure 1:

1. RF Front-End (unchanged)
    a. Passive signal conditioning and RF interfaces
    b. Insertion delays less than 10 ns, considered negligible in the overall latency budget
2. ADC/DAC subsystem
    a. Prototype high-speed converters featuring integrated digital signal processing (DDC, DUC, and EQ)
    b. Operates at double the sampling rate of legacy components (e.g., 800 MHz loop bandwidth)
    c. Significantly reduces ADC/DAC + DSP latency, measured at approximately 310 ns
3. Baseband DSP (FPGA-implemented)
    a. Redesigned RX and TX DSP chains using next-generation FPGAs with hardened DSP cores
    b. Incorporates a new ultra-low latency fading block with equalized routing paths across logic branches
    c. Measured latency of less than 50 ns for the fading block using timing-annotated simulation and waveform inspection
    d. Critical signal paths are moved to accelerated clock domains, improving throughput and deterministic behaviour
4. Data router and combiner block
    a. Architecturally optimized to reduce the number of serial bus transitions
    b. Transitioned from a three-hop to a single-hop routing structure
    c. Achieved ~150 ns latency, resulting in a net improvement of approximately 200 ns over the D5.2 design

To validate the latency contributions of individual subsystems, a dedicated loop test configuration was developed. This setup enables direct latency measurement across the ADC/DAC, DSP, and router blocks using both evaluation boards and integrated PoC implementations. These results serve as representative benchmarks for the final integrated system.

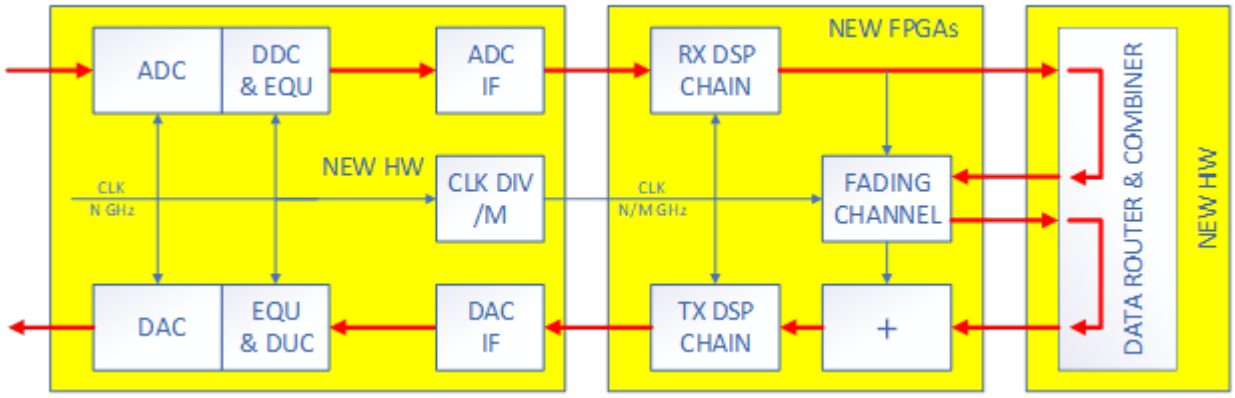

FIGURE 1. HIGH-LEVEL BLOCK DIAGRAM OF THE NEW POC ARCHITECTURE

2.1.3.2 ARCHITECTURE OVERVIEW AND FUNCTIONAL INNOVATIONS

The updated signal processing pipeline spans from analog RF signal capture to digital channel processing and output routing. Key architectural innovations include:

1. Integrated signal processing in converters
    a. The ADC and DAC modules now include embedded digital logic for DDC, DUC, and equalization
    b. Eliminates the need for these functions to be implemented in the FPGA, reducing data movement and processing latency
2. Clock domain optimization
    a. The addition of a clock divider/multiplier (CLK DIV/M) enables independent optimization of ADC/DAC and FPGA processing clocks
    b. Ensures synchronization while maximizing throughput in each subsystem's native domain
3. Enhanced DSP chain implementation
    a. Utilizes high-speed, low-latency hardened DSP slices within modern FPGAs
    b. Includes a fully re-architected fading channel block designed to maintain path-equalized latency across all routes
    c. Validated using timing waveform analysis with the shortest route measured at 32 ns and the longest at 41.6 ns
4. Latency-optimized routing and combining
    a. Supports multi-channel signal aggregation with significantly reduced serialization overhead
    b. Ensures timing coherence across MIMO paths and multi-emulator scenarios

2.1.3.3 HARDWARE INTEGRATION CONSIDERATIONS

The target platform consists of two modular Keysight hardware units, as it can be seen in Figure 2:

- A multi-port RF PROPSIM channel emulator chassis capable of scaling to 128 or more ports [4]
- A modular signal interface/controller unit providing synchronization and digital processing orchestration

FIGURE 2. KEYSIGHT PROPSIM F64 F8800B 128 RF PORTS CHANNEL EMULATOR

While the final D5.3 architecture is designed for integration into these platforms, vendor-side delays in component manufacturing (e.g., custom ADC/DAC devices) have impacted availability. As a result, the full system integration couldn't be completed before the end of the project timeline. Despite this, we were able to demonstrate what it was planned results-wise, achieving the defined project goals related to this PoC.

2.1.3.4 IMPLEMENTATION CONSTRAINTS AND MITIGATION

The following implementation constraints were identified and addressed:

1. Latency target: The architectural design achieves the target latency of less than or equal to 0.5 us, meeting 6G-SHINE in-X subnetwork D5.3 requirements.
2. Component availability: Prototype ADC/DAC modules remain in development due to external supplier delays. Current results are based on test boards and evaluation setups populated with early silicon and emulated clock conditions.
3. Proof-of-Concept Limitation: The final demonstrator will not be hosted in the production-grade hardware chassis. Instead, measured results are drawn from representative FPGA testbeds and converter loopback configurations, which are functionally equivalent in latency performance.

While the physical form factor may not reflect the final product, the measured latency, data flow architecture, and timing fidelity of this final PoC demonstrator are technically representative of the end product once integrated.

2.1.4 FINAL DEMONSTRATOR SCENARIO DESCRIPTION UPDATES/CHANGES

During the project timeline, it was agreed within WP5 that the final demonstrator scenario would incorporate RIS-assisted low-latency channel emulation. This section presents the outcome of that plan and showcases the emulation capabilities and performance evaluation of RIS-enabled propagation scenarios within the current PoC framework. The simulation scenario, illustrated in Figure 4 and Figure 5, models an indoor T-shaped corridor to evaluate RIS-assisted propagation under LOS-to-NLOS conditions. A transmitter (TX) is positioned at the base of the corridor, while a receiver (RX) moves along a predefined path that transitions from direct Line-of-Sight (LOS) to Non-Line-of-Sight (NLOS). A passive RIS is mounted at the intersection's apex wall to reflect signals into the NLOS region. The scenario is analysed at two frequency bands, FR1 (3.5 GHz) and FR3 (10 GHz), with a 100 MHz emulation bandwidth. The setup aims to characterize the RIS's effectiveness in improving coverage and channel quality by analysing parameters such as path loss, delay spread, and angular characteristics across the RX trajectory. In WP2 Deliverable D2.3 [1], we developed a comprehensive geometry-based RIS model designed for seamless integration with RF raytracing frameworks. In this work, the same RIS model [3], operating in focusing lens mode, where beams are symmetrically directed toward predefined focal points, was employed within CNIT's raytracing tool. Using this configuration, a T-shaped indoor corridor scenario was implemented to generate high-fidelity channel model data for RIS-assisted propagation analysis. This data was then imported into the Keysight GCM tool [4] to generate time-continuous channel emulation for use in the throughput evaluation test bed shown in Figure 3. A Keysight UXM 5G [5] network emulator provides the 5G signal, which is passed through the PROPSIM channel emulator to model realistic propagation, including RIS effects. A User Equipment (UE) device is placed in a shielded box to ensure a controlled RF environment. Performance is monitored using Keysight APT [6], which

captures throughput metrics for both with-RIS and without-RIS conditions. This setup enables accurate end-to-end testing of RIS performance impacts on link-level throughput in a reproducible laboratory environment.

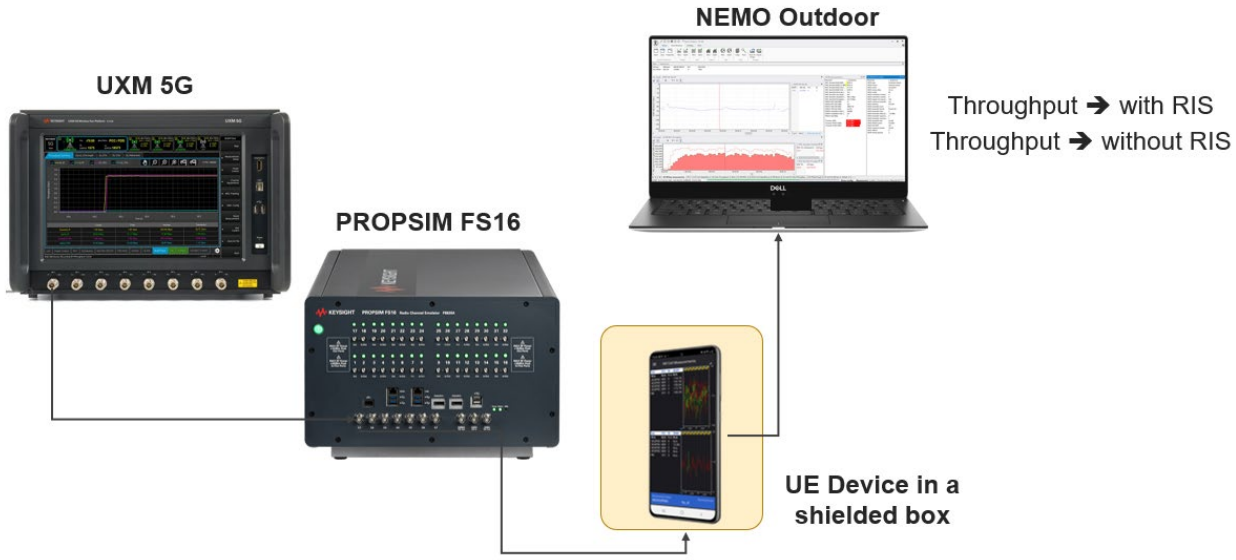

FIGURE 3. POC TESTBED SETUP FOR RIS-ASSISTED THROUGHPUT EVALUATION AT FR1

The RIS indoor 3D and 2D raytracing scenario implementation in GCM is shown in Figure 4 and Figure 5 respectively.

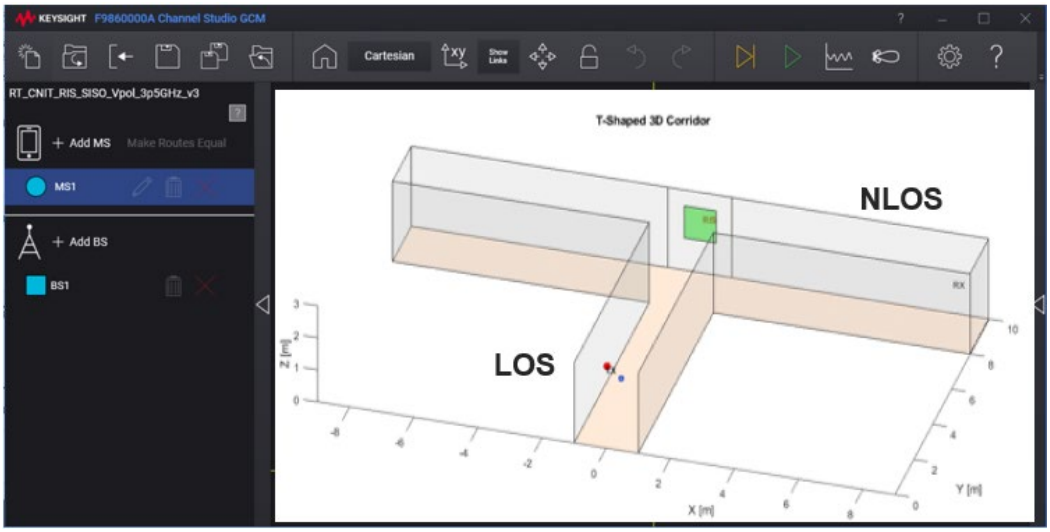

FIGURE 4. RIS INDOOR 3D RAYTRACING SCENARIO IN GCM (F9860000A CHANNEL STUDIO GCM, N.D.)

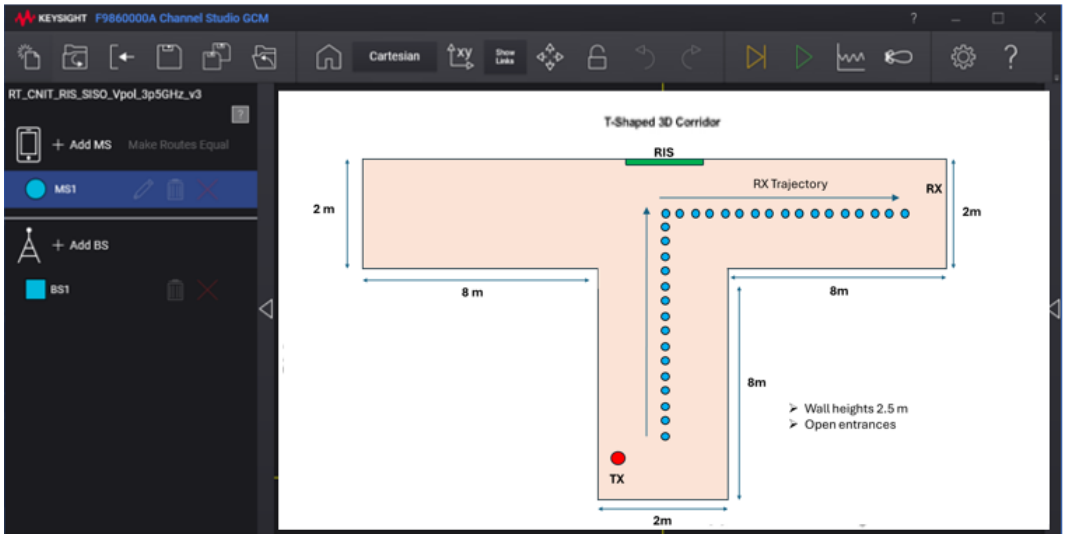

FIGURE 5. RIS INDOOR 2D RAYTRACING SCENARIO IN GCM (F9860000A CHANNEL STUDIO GCM, N.D.)

Figure 6 shows received power across the RX trajectory at FR1 3.5 GHz, comparing coverage with and without RIS in the indoor T-corridor scenario. In the LOS region, both curves are similar, but as the RX enters the NLOS section, received power drops sharply without RIS, frequently falling below the outage threshold (blue line). With RIS, signal strength remains consistently above the threshold, eliminating coverage holes. The outage probability is reduced from 20% (without RIS) to 0% (with RIS), demonstrating the RIS's effectiveness in mitigating NLOS degradation.

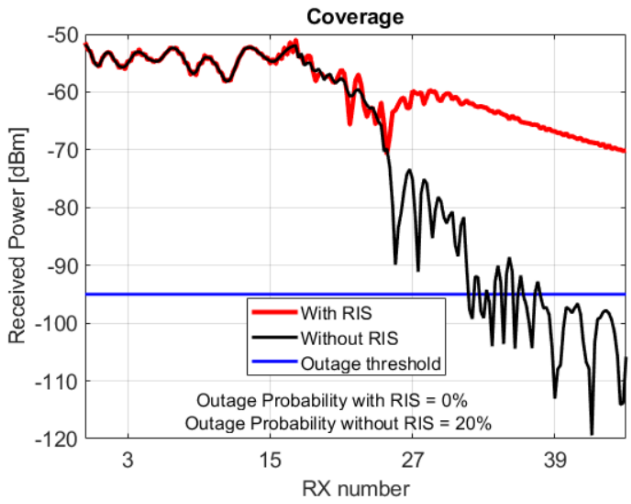

FIGURE 6. COVERAGE AT FR1 (3.5 GHZ) WITH AND WITHOUT RIS IN THE SCENARIO

Figure 7 shows the received power along the RX trajectory at FR3 10 GHz, comparing performance with and without RIS in the T-corridor scenario. In the LOS region, both cases perform similarly, but upon entering the NLOS region, signal degradation is severe without RIS, falling well below the outage threshold (blue line), resulting in a 42% outage probability. With RIS, the received power recovers significantly in NLOS, maintaining coverage and reducing the outage probability to just 0.1%. This highlights the critical role of RIS in mitigating high-frequency path loss and preserving connectivity in obstructed indoor environments.

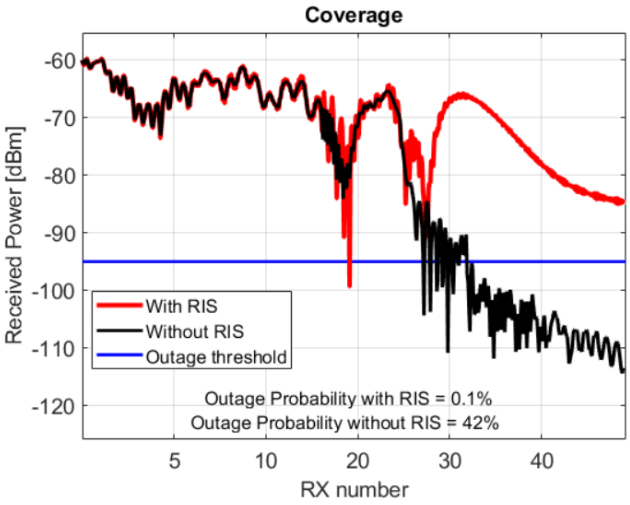

FIGURE 7. COVERAGE AT FR3 (10 GHZ) WITH AND WITHOUT RIS IN THE SCENARIO

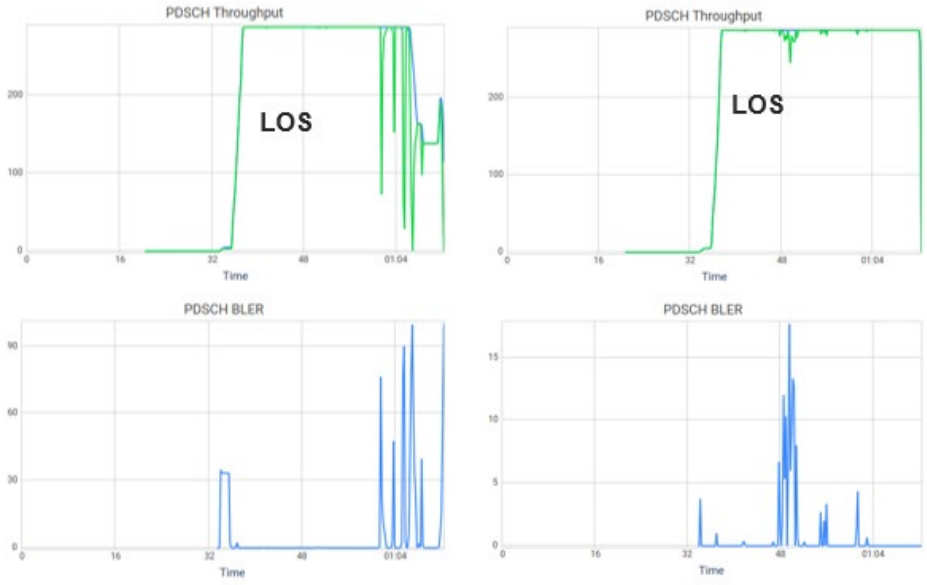

FIGURE 8. THROUGHPUT AND BLER EVALUATION FOR THE RIS-ASSISTED CHANNEL MODEL FOR THE INDOOR SCENARIO AT FR1

Figure 8 compares PDSCH throughput and BLER performance with and without RIS along an RX trajectory transitioning from LOS to NLOS:

- On the left (without RIS), throughput is high and stable in the LOS region but drops sharply in NLOS, accompanied by a significant increase in Block Error Rate (BLER). This indicates degraded link quality and unreliable data transmission under obstructed conditions.
- On the right (with RIS), the throughput remains high and nearly constant even in the NLOS region, while BLER is significantly reduced. These results confirm that the passive RIS effectively restores link quality, improving both throughput stability and error performance in challenging NLOS environments.

2.1.5 RESULTS

2.1.5.1 1.1.5.1 BASELINE RESULTS UPDATE

The baseline or reference latency for D5.2 was set to 2.1 µs. As presented in deliverable D5.1, the PoC achieved an insertion delay of approximately 0.9 µs through FPGA-level optimization of DSP blocks. That baseline serves as a reference for evaluating this second iteration of PoC implementation, which now incorporate hardware design improvements.

2.1.5.2 UPDATED TC RESULTS VS THE BASELINE AND/OR VS VERSION 1.0

The D5.2 baseline (~0.9 µs) remains valid for comparison. No significant updates to the original KPIs or measurement methodology were needed.

TABLE 2. COMPARISON OF LATENCY IMPROVEMENTS

Component	D5.2 Delay	D5.3 Delay	Improvement
ADC/DAC Loop	~600 ns	~300 ns	2×
Fading Block	~150 ns	~50 ns	3×
Router/Combiner	~300 ns	~100 ns	3×
Total Latency	**~0.9 µs**	**~0.45 µs**	**2×**

The qualitative gains are:

- Full signal chain response improved in jitter stability
- Faster rise/fall time under transient fading scenarios
- Synchronization performance validated across devices using multi-channel emulation

To ensure deterministic timing behaviour across the entire fading processing pipeline, it is critical that all signal paths, regardless of their internal routing complexity, experience equal and consistent latency. This guarantees temporal alignment of all outputs and prevents timing mismatches that could affect channel consistency or synchronization in time-critical communication scenarios. The figures below illustrate a detailed simulation analysis of the newly developed fading block architecture, where the goal was to evaluate and align the minimum and maximum latencies across all potential dataflow paths through the module.

Figure 9 captures the shortest possible route through the fading architecture. The observed minimum latency from input to output is 32.0 nanoseconds. This corresponds to the most direct path through the DSP pipeline, involving minimal combinatorial and routing logic. Figure 10 illustrates the longest internal route, with an observed latency of 41.6 nanoseconds. This extended delay results from traversing a more complex sequence of processing elements and internal bus transfers. In the figures, the green horizontal blocks represent the fading blocks, the red horizontal bar represents the streaming delay, the vertical white line represents the input trigger, and the vertical yellow line represents the output trigger.

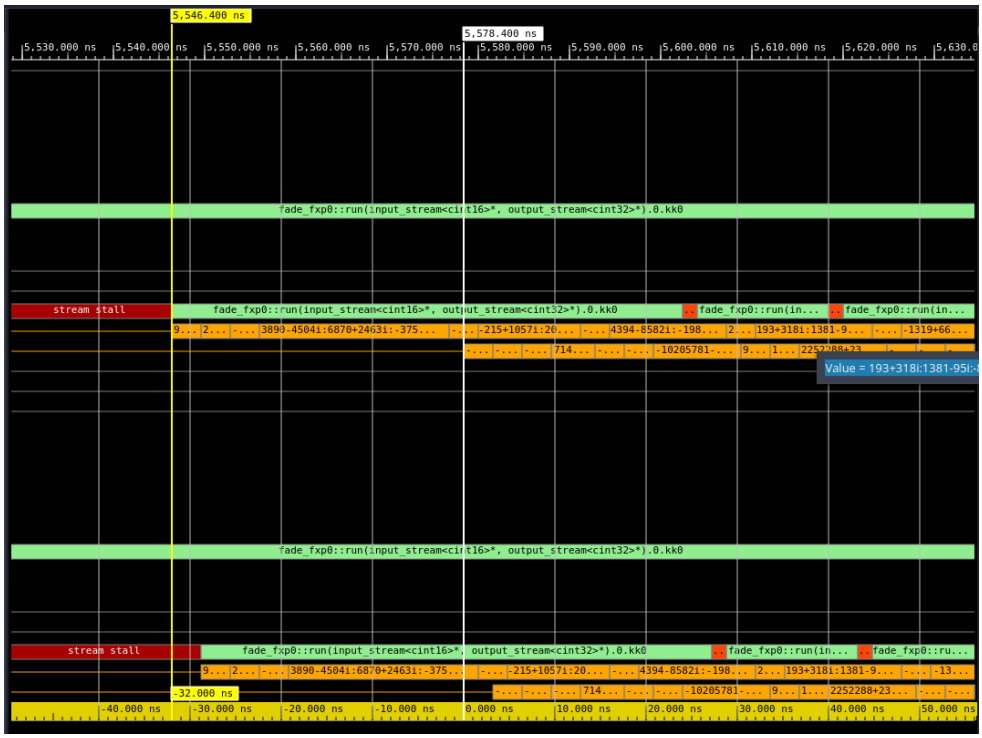

FIGURE 9. SIMULATION RESULTS FOR THE SHORTEST ROUTE THROUGH THE FADING ARCHITECTURE

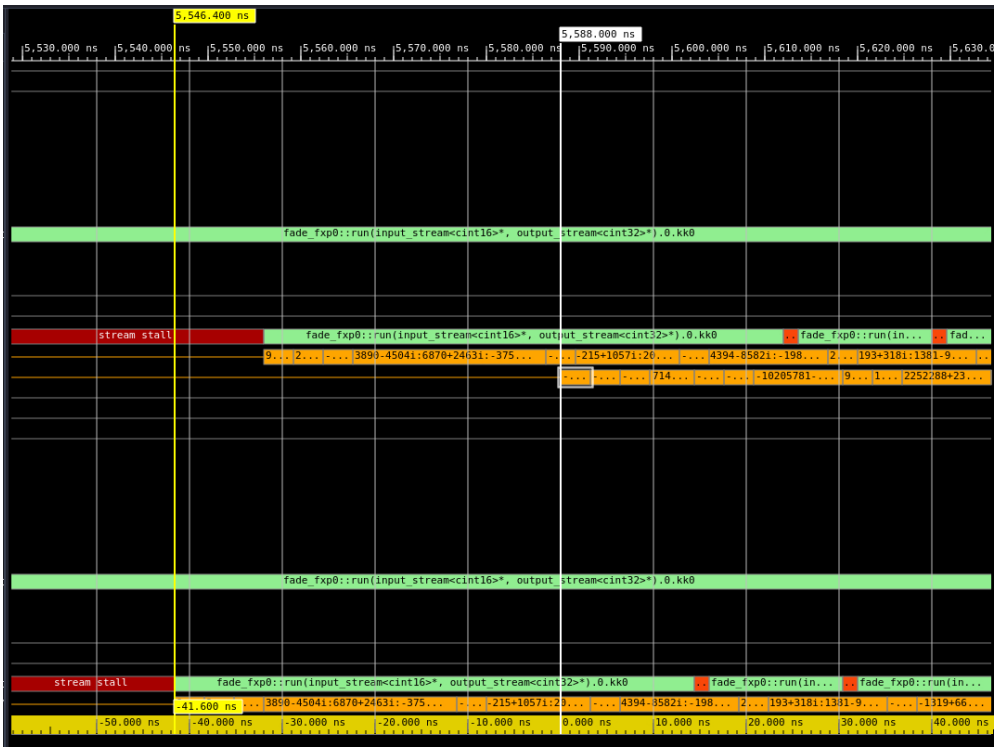

FIGURE 10. SIMULATION RESULTS FOR THE LONGEST ROUTE THROUGH THE FADING ARCHITECTURE

To guarantee uniform latency across all operational configurations, it is essential to compensate for path-length variations by introducing controlled pipeline delays on the shorter paths. Thus, all routes, irrespective of logic depth or functional complexity, must conform to the worst-case latency of 41.6 ns. This is typically achieved by inserting register stages or delay blocks at strategic points in the shortest paths.

Implementing this delay equalization ensures that:

1. All output samples emerge in lockstep, preserving timing alignment across channels.
2. The fading module operates deterministically, with constant latency.
3. It supports high throughput pipelining without risking race conditions or skew.

This analysis and tuning step is vital for meeting the ultra-low-latency performance targets (sub-50 ns for the fading block) in the D5.3 architecture of the emulator.

2.1.5.3 CONCLUSION

This PoC implementation demonstrates that carefully targeted enhancements in FPGA design, DSP pipeline optimization, and hardware architecture can successfully reduce radio channel emulation latency to 0.5 µs levels. The upgraded PoC implementation satisfies the channel emulator target set for 6G-SHINE in-X subnetwork latency

1. 2× reduction in total end-to-end latency
2. 3× performance improvement in critical DSP processing blocks
3. Seamless compatibility with next-generation ADC/DAC modules and modern FPGA platforms

These results establish a future-proof architectural baseline for advanced 6G testing applications, including closed-loop system validation, beam tracking, and ultra-reliable synchronization in short-range in-X scenarios. The key technical enablers include deep restructuring of the DSP signal path and integration of high-speed, low-latency converter technologies.

Overall, the project achieved a 77.3% latency improvement from the initial baseline at the project's inception to the final results reported in D5.3, representing a significant milestone in the evolution of low-latency channel emulation.

2.2 LATENCY AWARE MAC ACCESS FOR SHARED SPECTRUM BANDS

2.2.1 GENERAL DESCRIPTION

The proposed Latency Aware (LA) MAC TC was initially presented in WP3 in D3.1 and D3.3 deliverables [1] .We have now implemented the enhanced features presented in D3.3 also in the PoC and moved forward to measure its real life performance with the relevant KPIs.

This TC is focused on enabling the Robotic control use case and hence it is evaluated with relevant KPIs like jitter and latency (related to communication cycles) while adopting network traffic that models a typical robotic arm control traffic flow. By aiming to guarantee bounded maximum latency and jitter, the proposed approach enables subnetworks to support robotic control applications even in shared spectrum environments. This addresses one of the key limitations of current wireless communications: their inherently non-deterministic behaviour. While average latency and jitter may be acceptable and improve with increased physical-layer capacity, achieving strict bounds on these metrics remains infeasible -primarily due to limitations in the MAC access protocols used in shared spectrum scenarios, rather than physical-layer constraints.

In the following subchapters, the implementation of proof of concept of the version 2 (v2) Latency Aware MAC and the performance results acquired from testing under relevant conditions are presented.

For the sake of readability of this deliverable, we will continue with a short presentation of the proposed enhanced functionality of the LA MAC. For more details, we refer the reader to the extensive presentation of the LA v2 MAC in D3.3.

2.2.2 ENHANCED/NEW FUNCTIONALITY VS PREVIOUS VERSION

In the initial design of LA, the LA period access timing was fixed based on a predetermined interarrival time. It was found to struggle to support traffic flows that generate multiple packets per LAD period. The main issue identified is that if the total number of high priority packets generated per LA period is more than 2, then the LAD capacity is saturated, as there is no alternative way to forward the packets in time. As a result, packets were queued on the high priority queue, more nodes are trying to access each LAD period every time and finally more collisions are

probable in the LAD period (as more packets are competing during the LAD, the higher the chance of a LAD collision) and packets could end up waiting multiple LAD periods before being served. This had an impact on MAC access latency as well as packet success ratio, as more packets could reach the maximum retransmissions setting of 2 and get discarded.

To solve the aforementioned bottleneck and make the LA design able to cope with high priority traffic capacity issues in a scalable and abstract way, we redesigned LA to allow multiple LA packets to be strobed next to each other once an LA start time came. The proposed update to the LA design is presented in Figure 11 and is named LAD V2.0.

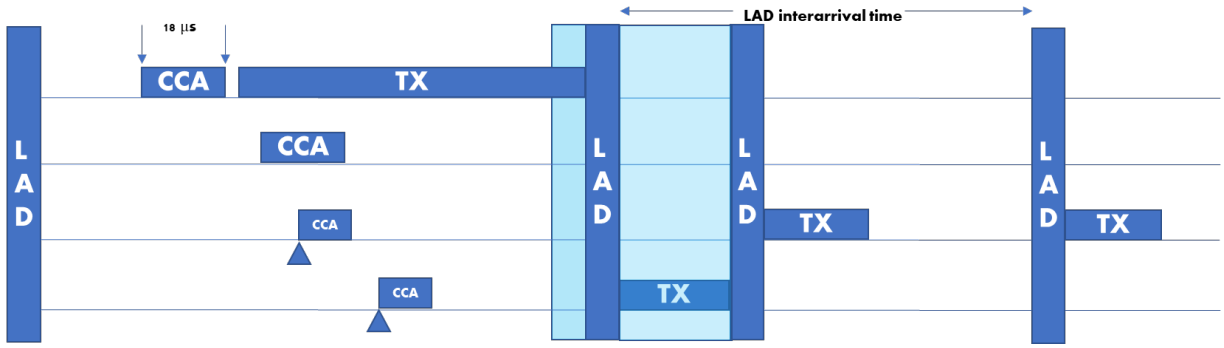

FIGURE 11. LAD V2.0 DESIGN

As it can be seen in the figure, when 2 or more nodes have a packet to send when an LA period starts, they are able to compete in the first LA, one will win and transmit its packet and immediately after the transmission of the first packet, a new LAD gets initiated from the nodes still having high priority traffic to send. In this way, a bottleneck is not created and in each LAD time, all high priority traffic is transmitted. The impact of this strobing approach on the legacy best effort access scheme was analysed in D3.3 and will also be evaluated using the PoC demonstrator.

2.2.3 HIGH LEVEL ARCHITECTURE CHANGES/UPDATES

The architecture of the implemented demonstrator was not changed from the previous implementation show in D5.2 [1]. It still is comprised of 2 LAD enabled devices that are based on a ZYNQ platform from AMD/Xilinx, namely the ZCU-102 development kit [17] equipped with Analog Devices FMCOMMS 2 RF frontend card [18], and 4 Intel

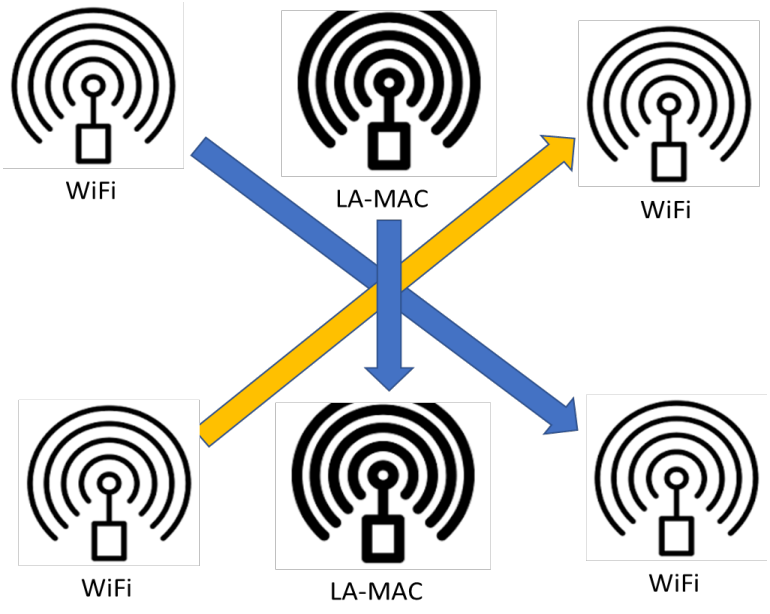

FIGURE 12. LAD-MAC DEMONSTRATOR TOPOLOGY AND STRUCTURE

NUC devices equipped with AX-200 Wi-Fi cards [19] that act as the ISM incumbent technology that compete with the proposed LAD equipped devices for access in the same channel.

The logical deployment layout is shown in Figure 12. All 6 nodes are placed close to each other in order to all reside in the same wireless collision domain. The transmission power of all systems is calibrated to deliver equal transmission power and avoid any PHY dominance effect. Using the implemented topology, traffic flows are injected through the active links and measurements based on Application and MAC layer KPIs are taken to determine the performance of the proposed LAD MAC, but also its impact on incumbent technologies.

2.2.4 FINAL DEMONSTRATOR SCENARIO DESCRIPTION UPDATES/CHANGES

Several changes were integrated into the demonstrator to support the v2 of the LA MAC. Mainly, the changes were on the software part of the MAC implementation, which we call the soft-MAC part, while the hard-MAC part, which resides in the FPGA and implements the TX FSM, remains the same. More specifically, the MAC scheduler was updated to be able to call for subsequent LA periods if more high priority packets are available in the high priority queue. Also, a maximum of consecutive LA packets was set to 50% of the time of the overall LA period in order to keep fairness and also allow lower priority traffic flows to have a chance to access the wireless medium. To force the packets to be sent always in a burst mode and not during the contention period, we disabled in this version of the demonstrator the ability of LAD to send packets also randomly in the non-LA period; hence the PoC is implementing now what is called pure LA version, as described in D3.3 and not the Hybrid-LA MAC version. This will surely impact the overall performance negatively, but it will allow us to accurately measure the burst feature impact on overall network performance which is the most important feature added in this final implementation of the PoC.

As shown in D2.2 [ref], a typical robotic link control flow requires a few Mbps, but it has very strict latency and jitter requirements in order for the link to be stable. To mimic a robotic arm control flow, we reused the flow design from D5.2. We choose 2 flows (one in LAD link and one in Wi-Fi_P link) to be a representation of a typical 2 ms control loop implemented in a PLC-slave robot industrial setup. In Table 3 the settings of the active flows in the demonstrator are presented.

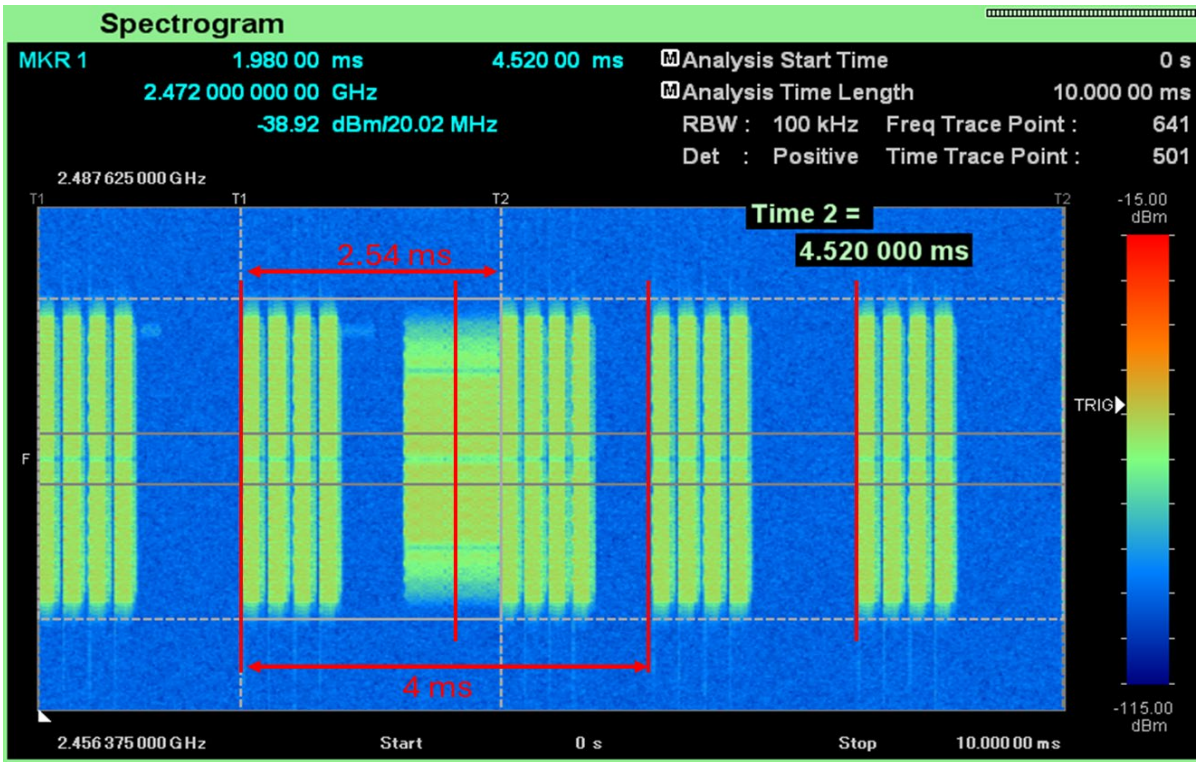

FIGURE 13. BURST FEATURE TEST AND DELAYED LA PERIOD CORRECT OPERATION VALIDATION

TABLE 3. TRAFFIC FLOW SETUP

LAD robotic arm control link		Wi-Fi 1 robotic arm control link		Wi-Fi2 background traffic	
PPS	Packet size	PPS	Packet size	PPS	Packet size
500	300 bytes	500	300 bytes	500-2500	1400
1.2 Mbps		1.2 Mbps		Varying from 5.6 to 28 Mbps	

While the specific flow profiles remain the same, we now inject more than one high priority flow in each node, in order to test the burst packet feature of the v2 LA MAC compared to how Wi-Fi would handle multiple high priority flows.

In order to test the correct operation of the implemented v2 LA, we executed various test scenarios where multiple high priority traffic flows were injected in the LA nodes and observed the results in real time using a spectrum analyser, but also by looking at the calculated KPIs per experiment.

In Figure 13, we can see what the spectrum looks like when 4 robotic arm control traffic flows are injected in the network from the same node, when configured to use the same 2 ms packet interarrival time and are also synchronized in time.

This traffic flow setup results in 4 packets becoming available every 2 ms in the same node and need to be sent as soon as possible. By observing the captured snapshot, we can see how LA handles the presence of a 3rd party technology packet in the air. By following the LA Sense and Wait approach and not immediately backing off when the medium is captured, LA MAC is able to wait for an active transmission to finish, and then by waiting for a random time from 10-30 μs το deconflict with other LA enabled nodes it is able to get access to the medium with minimum latency. Since burst is supported in LA v2 MAC, all 4 packets are sent once the medium is captured after the external packet has finalized its transmission. As it can be seen, the periodic LA start time was scheduled to be initiated during the external packet transmission period, resulting in an added latency of 2.54 – 2 = 0.54 msec and relevant jitter due to waiting for the external packet transmission to be finalized. In any case, the offered latency

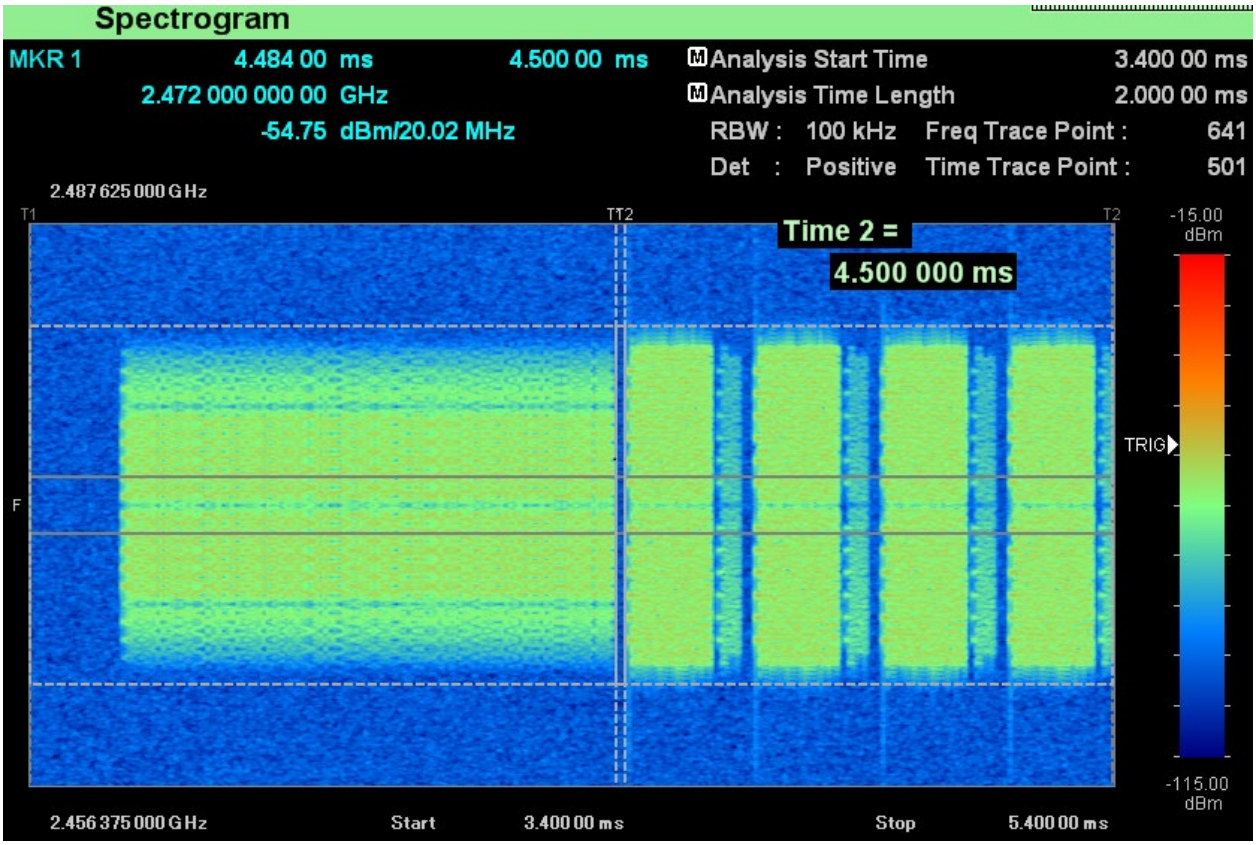

FIGURE 14. SENSE AND WAIT LA MAC ACCESS IN ACTION

is significantly better than what LBT with a backoff supported approach would achieve, where the packets would be waiting for a random backoff timer to expire to try to get access to the medium again. It is also important to note that the next cluster of packets that came in the next 2 ms period were treated at the correct timing, as LA is not cumulative in the induced delay from the previous delayed LA access. Hence, the next cluster of packets was sent exactly 4 ms after the one before the external packet was present, as should be if we take into account the 2 ms LA period used in this experiment.

For this specific snapshot, LA was able to get access to the medium only 16 µs after the external packet transmission was concluded as can be seen in Figure 14. The efficiency of the Sense and Wait approach, combined with the burst transmission feature, helped minimize the impact of external traffic on the medium. As a result, the 4 priority packets waiting at the LA node experienced reduced medium access delays, as well as lower overall latency and jitter at the application layer.

To evaluate the burst traffic feature, we set up a series of experiments with a different setup of traffic flows from the previous version of the PoC. Now, each node services 2 flows of priority traffic while the background traffic remains a single flow of periodic traffic. The 2 priority traffic flows in each link were set to create a packet each at the exact same time, forcing the wireless stack to service 2 packets arriving almost instantaneously as soon as possible. The overall traffic load of the network is depicted in Table 4 for all 6 experiments vs background traffic flow. The background traffic ranged from 0 to 2000 packets per second with a step of 400 packets per second, hence 6 experiments in total. Each experiment was executed for 200 seconds and results were gathered from each node for the running time.

TABLE 4. EXPERIMENT SERIES CONFIGURATION

		Experiment #		1	2	3	4	5	6
	App Pkt size (bytes)	Airtime (µs)							
Background	1400	298	**PPS**	0	400	800	1200	1600	2000
Wi-Fi priority	250	123		500	500	500	500	500	500
LAD priority	250	216		500	500	500	500	500	500
Channel utilization				17.0%	28.9%	40.8%	52.7%	64.6%	76.6%

The TX power of the Wi-Fi nodes was adjusted to match the TX power of the SDR based LA implementation, while the topology of the testbed was arranged so all transmitter nodes where close to each other so they can sense at the same power level the transmissions of each other.

2.2.5 RESULTS

2.2.5.1 BASELINE RESULTS UPDATE

The baseline for this experiment is updated as the incoming traffic is now altered. We made a direct comparison of what a typical ETSI complaint device can do using CSMA/CA against what the proposed LA PoC can offer in the same networking and wireless conditions (same load, came collision domain). Hence, we will focus comparing the LA traffic performance and the Wi-Fi priority flow performance in terms of Jitter, latency and throughput stability as the background traffic changes. We should point out at this time that the Wi-Fi nodes support a higher PHY rate up to 48 Mbps compared to the SDR based LAD implementation that supports only up to 11 Mbps, hence the almost half airtime of a Wi-Fi priority packet compared to the same size in bytes LA packet. But even with that difference in actual airtime of a single packet, the results can be revealing of the performance gains of the LA TC.

2.2.5.2 UPDATED POC RESULTS AND COMPARISON VS WP3 THEORETICAL RESULTS

Starting the result presentation, we will focus on the general performance achieved from LA vs the Wi-Fi priority link while background traffic was ranging from 0 to ~75% channel capacity. We will then try to explain and correlate these results with what was observed in the v1.0 of this PoC demonstrator and the relevant features of the theoretical design.

In Figure 15, we can see the achieved latency of both priority and background traffic flows. A discussion on what we can observe in this figure is surely needed. Each experiment is enclosed in a black rectangular so it's easier for the viewer to realize which flows belong to the same experiment. Starting from the left where experiment 0 is presented with no background traffic, we move all the way to the right where the last experiment, experiment 6, is presented with maximum background traffic.

Some initial observations are that LA access scheme seems indeed to manage quite easily to support multiple flows, offering very stable latency in all traffic load scenarios. From 0 traffic to max traffic in the last experiment, LA MAC only suffers a 17% increase in each average time, in absolute values that is from 1.09 ms to 1.28 ms, an average latency increase of just 190 µs. Regarding max, min and jitter of latency the numbers are also revealing, only 2% higher min latency, 3% higher max latency and 56% higher jitter. That is, from operating in parallel with competing traffic up to a 75% saturated wireless channel.

On the opposite side the Wi-Fi priority link suffered a devastating 6678% increase in average latency, in absolute numbers from 1.07 ms it went up to 71.25 ms average latency. Minimum latency also grew by 19% while maximum latency grew by 709%. Jitter for the Wi-Fi priority links was also significantly deteriorated, increasing by 4607%. Based on these numbers but also on the general presentation of the graph in Figure 15, we can safely say that the proposed LA MAC is offering a quite stable and near-deterministic latency-related service. Load of the channel has little to no impact on the performance of the LA MAC which is something unique for spectrum sharing MAC schemes. There is no other MAC design that we are aware of, that can offer stable QoS in shared spectrum bands when sharing the spectrum with heterogeneous technologies up to date.

Comparing directly LA and Wi-Fi priority latency, we see that LA MAC was able to offer equal performance for low network loads, but up to 98% lower latency (93% was achieved in simulations) for high network loads, while maximum latency and jitter also where improved by an astounding 98.6% and 99.6% respectively (95% and 90% in simulations) in the worst network load condition scenarios. These results are fully in line with what was observed in D3.3 for the burst mode analysis using simulation tools where we observed also a very stable operation in much higher density scaled deployments than the 6 nodes we were able to deploy in the PoC demonstrator. Another very important observation is that LA MAC offers a bounded latency in all cases, no more than 8.72 msec in worst case during the full extent of the experiment series. We would like to explain why that is possible and to what extent we can talk about a deterministic latency TC. As LA MAC is able to get priority access to the medium, then the main delay factor that contributes to the max latency it can exhibit is the presence of other technologies that capture the medium and keep it for a long time. Worst case scenario analysis is that another wireless technology gets access to the medium and sends a very long packet that can last for a certain time. LA will need to wait for the packet to be transmitted, hence, according to worst case scenario analysis, the LA MAC will exhibit a delay equal to the LA period plus the packet transmission duration by the external technology. As we fully control LA period and we can set it to a fixed value, like 1 or 2 ms, the main uncontrollable latency contribution factor is the external packet duration. If some technology captures the medium in shared bands, regulation offers various rules to set a limit on Channel Occupancy Time (COT). ETSI for instance defined COT as a maximum of 60 ms for frame-based equipment and 13 ms for Load based equipment in [20] . Hence the bounded latency that LA MAC can offer depends on the current maximum allowed packet size in ISM bands and can be up to 60 ms. We would argue that ETSI should also update this max limit for industrial uses and applications so that to limit the max packet size into much smaller packets, if possible, below 10 ms in order to allow a bounded latency offering of LA below the 10 ms mark in any case.

We should also look at the throughput results to determine if the stable operation of LA was also maintained as well under high traffic loads and what was the packet loss in these cases. In Figure 16 the full profile of the resulting traffic flow throughput is presented for each experiment. We chose to omit the background traffic data points as

they are not so relevant for this comparison and they would scale out the graph, not allowing us to see the subtle but present differences between the similar robot arm control links serviced by LA MAC and Wi-Fi. The experiments presented in the graph are from left to right, low background traffic or experiment 0, to highest background traffic or experiment 6 as it was presented in Table 4. As we can see also in the throughput graph, the LA MAC offers quite a stable throughput in all load conditions. When compared with the similar Wi-Fi robot control links, we can see that the deterioration of QoS in the Wi-Fi server links is evident also as far as throughput is concerned while the LA MAC throughput shows very little differentiation between 0 and max background traffic load experiments.

The deterioration of the LA links is limited to 1.58% for average throughput vs 3.81% for the Wi-Fi priority flows from low to high traffic load scenarios. What changes rapidly though is the max and minimum throughput values achieved from the Wi-Fi links while the LA MAC links remains mainly stable. Specifically, LA MAC only loses 3.76% minimum throughput, and 0.76% higher max throughput compared to the 40.02% lower minimum throughput and 12.81% higher max throughput for the Wi-Fi priority links. The variation of min and max of the Wi-Fi links shows clearly the instability of the connection under load, while LA MAC is not significantly affected by higher background network load. Such a result marks on more tick towards deterministic networking, not only supporting bounded latency but also very low packet loss due to collisions, that in return also minimized the need for packet retransmissions and hence latency.

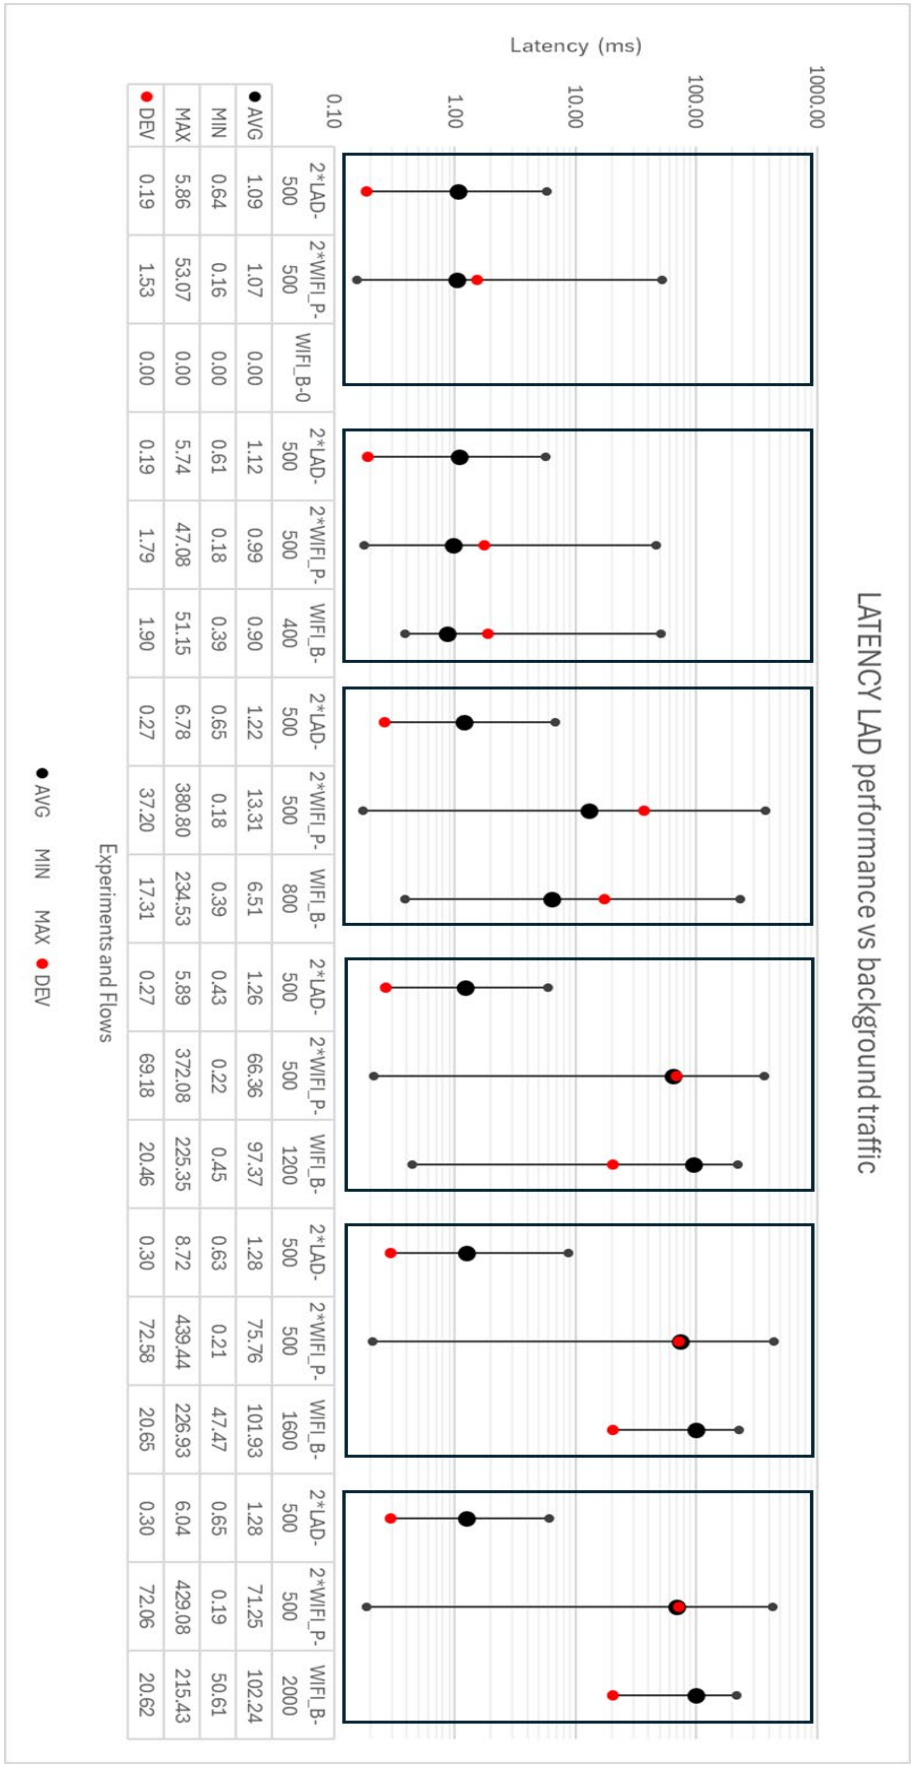

	2*LAD-500 2*WIFI_P-500 WIFI_B-0	2*LAD-500 2*WIFI_P-500 WIFI_B-400	2*LAD-500 2*WIFI_P-500 WIFI_B-800	2*LAD-500 2*WIFI_P-500 WIFI_B-1200	2*LAD-500 2*WIFI_P-500 WIFI_B-1600	2*LAD-500 2*WIFI_P-500 WIFI_B-2000
AVG	1.09	1.12	1.22	1.26	1.28	1.28
MIN	0.64	0.61	0.65	0.43	0.63	0.65
MAX	5.86	5.74	6.78	5.89	8.72	6.04
DEV	0.19	0.19	0.27	0.27	0.30	0.30
	2*WIFI_P-500 WIFI_B-0	2*WIFI_P-500 WIFI_B-400	2*WIFI_P-500 WIFI_B-800	2*WIFI_P-500 WIFI_B-1200	2*WIFI_P-500 WIFI_B-1600	2*WIFI_P-500 WIFI_B-2000
AVG	1.07	0.99	13.31	66.36	75.76	71.25
MIN	0.16	0.18	0.18	0.22	0.21	0.19
MAX	53.07	47.08	380.80	372.08	439.44	429.08
DEV	1.53	1.79	37.20	69.18	72.58	72.06
	WIFI_B-0	WIFI_B-400	WIFI_B-800	WIFI_B-1200	WIFI_B-1600	WIFI_B-2000
AVG	0.00	0.90	6.51	97.37	101.93	102.24
MIN	0.00	0.39	0.39	0.45	47.47	50.61
MAX	0.00	51.15	234.53	225.35	226.93	215.43
DEV	0.00	1.90	17.31	20.46	20.65	20.62

FIGURE 15. LATENCY RESULTS FOR LA AND WI-FI ROBOTIC ARM CONTROL LINKS VS BACKGROUND TRAFFIC

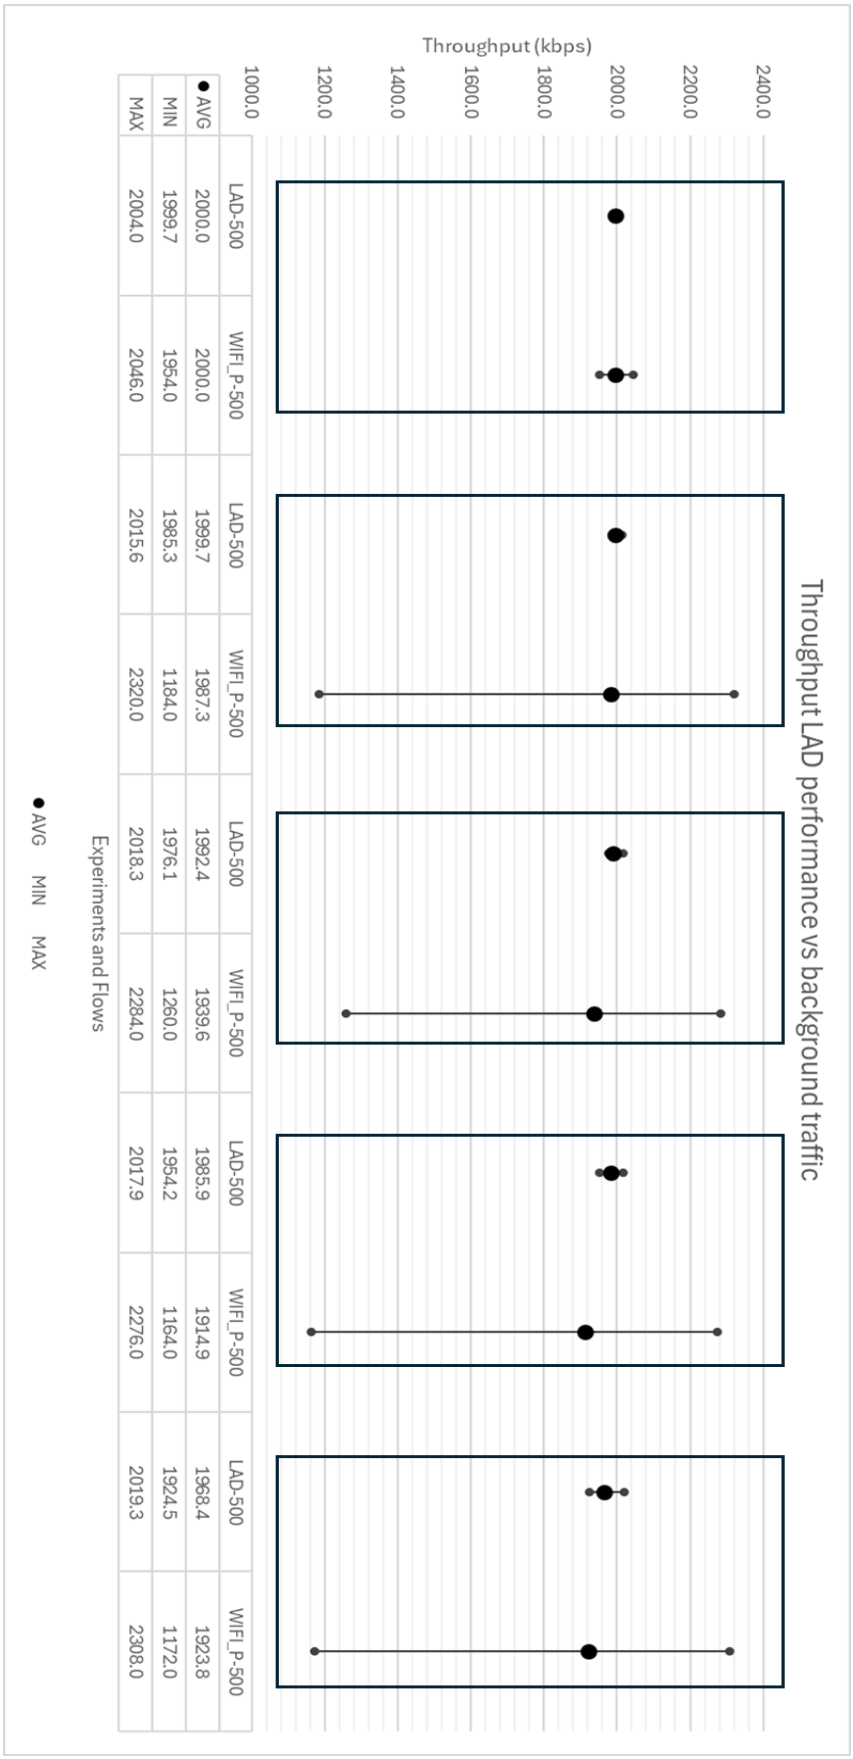

FIGURE 16. THROUGHPUT RESULTS VS NETWORK LOAD

To see in more detail the achieved latency performance, we present in Figure 17 the CDF of latency for both the LA and Wi-Fi served robotic arm control flows. We do not present all experiment graphs, as based on the overall performance we could identify 3 distinct cases, one where Wi-Fi links are still stable, a transient phase where performance starts to deteriorate and maximum is increased, and one where the Wi-Fi links are now saturated and suffer greatly from the increased load of the network.

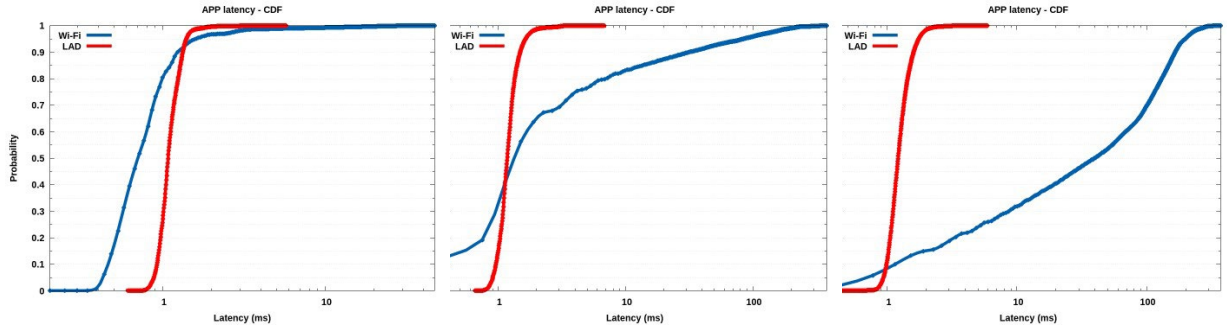

FIGURE 17. STABLE, TRANSIENT AND UNSTABLE PHASE OF THE WI-FI LINKS (A,B,C) RELATED TO OFFERED LATENCY

We observed that experiments 1-2 results are similar to Figure 17-A, experiments 3-4 are very similar to Figure 17-B and finally the last 2 experiments where like Figure 17-C. This clearly shows the gradual deterioration of the Wi-Fi served priority links while again it is evident that the LA MAC offers a very stable and deterministic operation, not getting impacted from added background traffic.

A similar picture is presented in Figure 18, regarding application layer jitters specifically between all 6 experiments in this experimentation series.

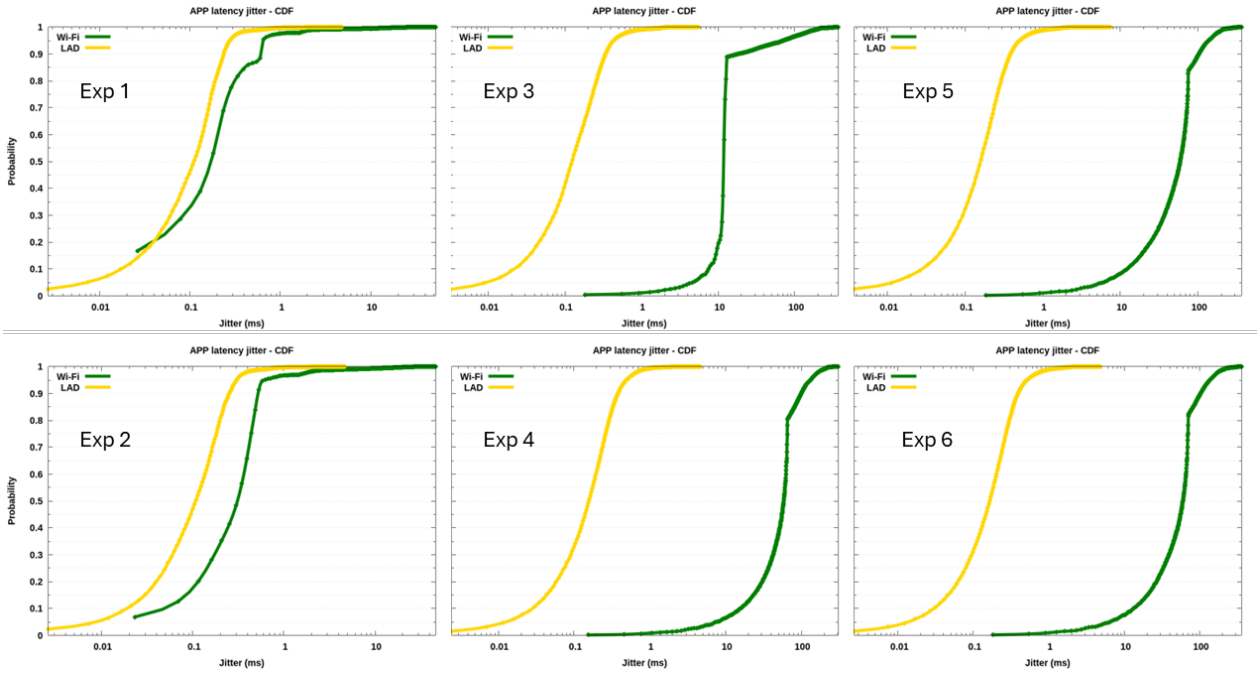

FIGURE 18. APP JITTER CDF FOR ALL EXPERIMENTS , WI-FI VS LAD FLOWS

It is evident that in all cases LA jitter is lower than the achieved jitter in the Wi-Fi served traffic flows. But again, as the network load increases, the Wi-Fi based flow's jitter significantly deteriorates from similar at low network loads (experiments 1 and 2) all the way to 100 times higher jitter than the LAD MAC offered one (experiments 4-5-6). As jitter is an important KPI for robotic arm control applications, lowering the jitter using the LA MAC TC over 100 times compared to a typical ISM wireless technology like Wi-Fi clearly shows that LA MAC can be a very significant enabler for wireless deterministic communications.

2.2.5.3 CONCLUSION

The evaluation of the PoC performance and its comparison vs the defined baseline showed expected gains in terms of latency and jitter to be similar in very low network load scenarios but as much as 98.6-99.6% lower respectively in high network load scenarios, which is a tremendous improvement. Interestingly, the impact of higher network loads to the LA MAC served flows was hardly visible, leading to the conclusion that LA MAC can indeed offer true hard prioritization against best effort access practices employed today in ISM bands and can offer a very stable performance irrelevant of network load. These results are in line with the simulation studies presented in D3.3 even though the scenarios simulated were not identical in density. In the simulation we went from 8 all the way to 32 nodes competing in the same collision domain while in the PoC demonstrator we were able to have available only 2 LAD enabled nodes and 4 Wi-Fi nodes. But still that was enough to identify the convergence and same trend of behaviour qualitatively between simulation and PoC.

Such a significant improvement in both latency and jitter is surely paving the way towards the realization of industrial robotic arm control applications in subnetworks, as long as they are using wireless systems that can support the LA MAC TC. The support for multiple hard prioritized traffic flows and the minimization of latency by supporting the proposed burst feature of LA MAC solves one of the main issues identified in D3.1 regarding the periodical LA proposed scheme, that is the inability to treat more than one packets per LAD period in the pure LAD design. The other future feature that can help is the Hybrid LA feature, where priority traffic is able to get access to the medium on-demand also during the non-LA period; however, this feature was not implemented in this PoC and it is left as a next step for future research. Even with this missing feature in the PoC, the burst feature is proven to be enough to support multiple packets per LA period and offer still stable and near deterministic wireless networking.

We would like to evaluate also the hybrid LA MAC scheme in our demonstrator in the near future and if possible support also higher density network deployments in order to verify hybrid LA MAC's impact in latency and reliability as well as prove the resilience of the proposed MAC design in higher density situation, as it was already indicated by the simulation based evaluation in WP3.

2.3 JAMMER RESILIENT PHY: WI-FI6 BASED CTI DETECTION AND MITIGATION BY OFDMA

2.3.1 GENERAL DESCRIPTION

In D3.3, we presented a method to detect and identify jammer type, and briefly showcased how it could improve scheduling in an Wi-Fi compliant OFDMA network. In this deliverable, we extend and implement the interference mitigation techniques and prototyped it on the openwifi platform. In comparison to the simple measurement of D3.3, the evaluation is now more complete. For instance, the jammer's duty cycle is kept at a fixed value in D3.3, here it varies between 0% to 100%. Also, in comparison with D3.3 there are more jamming mitigation techniques reported in this deliverable. For instance, in D3.3 only "reallocating the resource unit" is briefly tested, while in this deliverable "puncturing resource unit" and its real-time analysis on physical layer are also included. This provides direct evidence that the chances of consecutive collision will be significantly reduced, and link stability and latency will be improved. As a result, this technology component benefits the use case "Robot control" defined in WP2. Furthermore, by applying the technology, we can sustain same link quality despite the presence of the jammer, the performance loss is kept very low.

The real-life experimental evaluation results are detailed in the subsequent sections, for more complete analysis, including relevant simulations and theory, please refer to [16].

2.3.2 ENHANCED/NEW FUNCTIONALITY VS PREVIOUS VERSION

In D5.2 we focused on improving jammer resilience on the physical layer for a wideband jammer using an approximated log-likelihood based demapper. In D5.3 the prototype is enhanced to also cope with narrowband jammer by performing interference classification across heterogeneous technologies and applying appropriate

mitigation techniques. Hence the narrow band jammer is also referred to as Cross-Technology Interference (CTI) in the subsequent sections.

Multi-user OFDMA (MU-OFDMA) is one of the key features in Wi-Fi and cellular technologies, which will certainly be employed in 6G networks. While the Wi-Fi standard does define the packet format and waveform clearly, it did not specify how and when to apply, leaving a vast range of possibilities. More specifically, in earlier Wi-Fi standards, a packet to a user always occupies full channel bandwidth, but since Wi-Fi6, a user's data can occupy only part of the total bandwidth, which is referred to as the Resource Unit (RU). The method to allocate the RUs to users is purely up to chip vendors. Therefore, we research how to optimally perform RU allocation in order to achieve best system performance.

What is also worth mentioning is that since this PoC is complementary to the previous PoC in D5.2, in the sense that they improve robustness against different types of jammers (i.e., wideband vs narrow band), the baseline of this PoC is thus not the results obtained in earlier deliverables. Instead, we choose to use the performance when no OFDMA scheduling is performed as the baseline. With the standard compliant terminology, this is often referred as result obtained using Single User (SU) packet format.

2.3.3 JAMMER MITIGATION TECHNIQUES AND HIGH-LEVEL DEMO SCENARIO

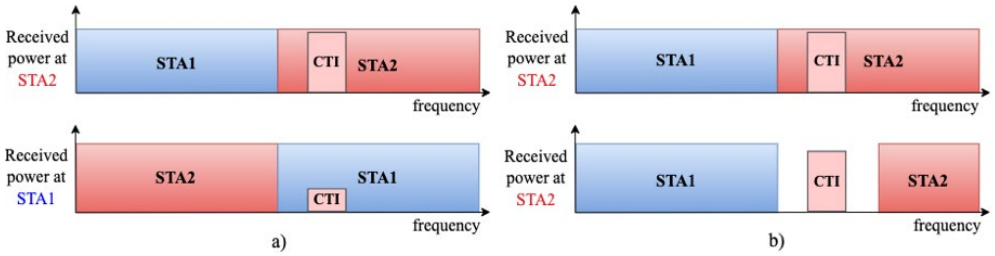

FIGURE 19. MITIGATION TECHNIQUES IN THE PROTOTYPE: A) RU REALLOCATION, B) RU PUNCTURE

Among several mitigation techniques, two specific ones are selected for prototyping. The first one is called RU reallocation. Since stations can be in different locations surrounding the AP, some might experience more CTI than others, especially when the interferer's transmit power is low. Then, allocating an unaffected RU to the user that experiences the most CTI will limit the impact of the interference. This is shown in Figure 19 by the received power at the STA2 before, and at STA1 after RU reallocation. The biggest advantage of reallocating RUs is that it comes at no cost when there is no CTI, since switching RU indexes of the same width does not influence throughput, assuming equal channel conditions. However, it will only be effective for downlink packets when there is a significant difference in the level of CTI for the different stations. Furthermore, it is only suitable when there is an unaffected RU available with a bandwidth suited to satisfy the users' traffic requirements.

The next option is to leave an affected RU unoccupied, i.e. *RU puncturing*. For downlink, this is allowed as long as at least 4x26 subcarriers per 20MHz bandwidth are occupied. For uplink, there is no lower bound for the number of occupied subcarriers. When the overall transmit power remains equal to the case without punctured RUs, the average power on the remaining RUs is higher when decreasing N_{RU}, as indicated by the equation below. The key principle of this equation is to keep the total transmit power the same as if no RU is punctured. For more concrete information about this equation please refer to IEEE 802.11ax standard, Eq 27-3 [7]. This is allowed as long as the regulatory limit for the power spectral density is not exceeded, which is 10 dBm/MHz for ETSI [7] and 17 dBm/MHz for the FCC [9].

EQUATION 1. SCALING FACTOR OF POWER PER RU

$$s_r = \frac{\alpha_r}{\sqrt{\sum_{r'=0}^{N_{\mathrm{RU}}-1} \alpha_{r'}^2 |K_{r'}|}},$$

where α_r is the power boost factor at the RU, N_{RU} is the number of occupied RUs, and K_r is set of modulated subcarriers in the RU with index r.

RU puncturing is mostly suitable for a relatively strong CTI, or when the impact of Wi-Fi on the other technology should be minimized. Another benefit of this method is that it allows Wi-Fi to estimate the Received Signal Strength (RSS) of the interference using CSI, since this portion of the spectrum is then not occupied by Wi-Fi signals.

After introducing the mitigation techniques, we discuss the methodology for CTI-aware OFDMA scheduling used in this PoC. In general, the same mechanism can be applied for uplink and downlink CTI-aware scheduling, except that for downlink the AP needs to be aware of the CTI status at each station, whereas for uplink the AP only requires to know its own CTI status, and not that of each receiver. Another minor difference is that for downlink the AP directly assigns various configurations in its data packets, while for uplink, the AP needs to control the OFDMA parameters via trigger frames as specified in the 802.11ax standard.

In order to detect interference and choose the appropriate mitigation technique, we propose to combine CTI feedback from STAs with a per-RU CCA at the AP. The methodology is explained hereafter and schematically visualized in Figure 20 with one AP and two STAs.

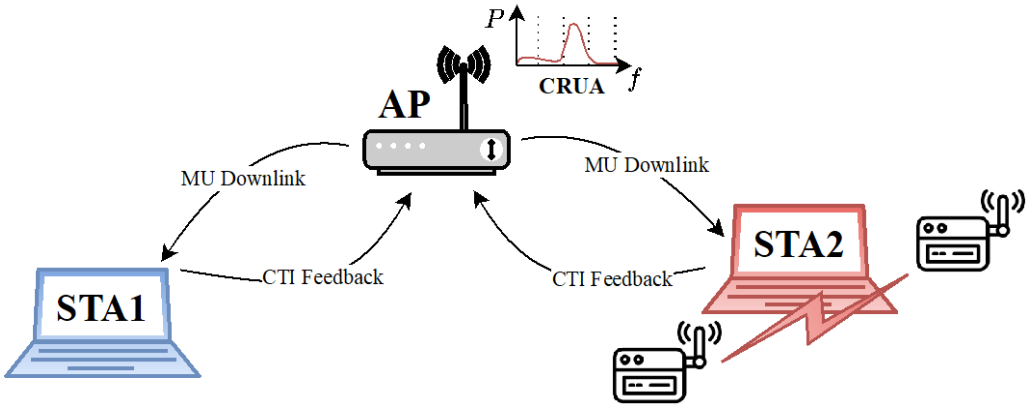

FIGURE 20. VISUALIZATION OF THE METHODOLOGY FOR CTI-AWARE OFDMA SCHEDULING.

To ensure that the level of CTI at the receiver is measured, STAs will perform the detection themselves and report the result to the AP, including the spectral location. We have shown in [10] that the detection can be done using the CSI from regular received Wi-Fi packets, thus requiring no hardware modifications or additional overhead in the form of control packets or interference mitigation pilot subcarriers. Note that for efficient uplink OFDMA scheduling, STAs already need to send buffer status reports to the AP. Furthermore, they may send bandwidth query reports that include the busy or idle status per 20 MHz bandwidth segment. CTI classification is a similar indicator, and the result of several CTI detections can be piggy-bagged in these status report packets to avoid too much overhead.

There are some downsides to CTI feedback from the STAs based on CSI. Firstly, if CSI is only reported for detected packets, a strong interferer may cause a packet to become undetectable, thus preventing the interference detection based on CSI to even start. To overcome this, the AP itself can also record the Packet Error Rate (PER) statistics per RU when it is alternating RUs. Furthermore, CTI feedback cannot be done in real-time as it requires receiving a packet and transmitting the feedback to the sender. Especially for technologies like Bluetooth that use frequency hopping, CTI feedback might not be a good indicator for any potential future CTI presence. Therefore, we propose to perform real-time interference detection at the AP as well, as explained in the following paragraph.

Since each Wi-Fi receiver will have a Fast Fourier Transform (FFT) module for demodulation, it could also be used to perform a spectrum scan in the working channel during an idle period. In this way, the AP can detect an ongoing interference packet, if there is sufficient CTI power at the AP. In essence, this is a CCA per RU, or Clear Resource Unit Assessment (CRUA) as we will call it from now on. It is similar to, but simpler than the enhanced CCA proposed in [11] and [12]. Just before preparing a packet transmission, CRUA will determine whether one of the RUs experiences interference by measuring its received power relative to the other RUs. For example, RU index 1 will be judged as interfered when the following statement holds:

EQUATION 2. CURA CTI DETECTION CONDITION

$$\tau \cdot P_{\mathrm{RU}_1} > \frac{1}{N_{\mathrm{RU}} - 1} \sum_{i=2}^{N_{\mathrm{RU}}} P_{\mathrm{RU}_i}, \quad 0 < \tau < 1,$$

where τ is a threshold value, and P_{RU_i} is the average power of RU index *i*. The threshold value τ is typically set to at most 0.5 to avoid false positives and can be adjusted if STAs report interference at a certain RU already. Note that devices will defer transmitting when the legacy CCA procedure considers the channel busy, although CTI is often not detected due to the difference in bandwidth and the relatively high CCA threshold.

Using the aforementioned methods, the AP can now make an informed decision for the OFDMA allocation. Either using direct feedback from the stations or based on higher-level statistics from rate control algorithms, an estimation for the SINR not only per user, but also per RU can be determined. Note that this results into a high-dimensional problem, due to the following aspects:

- There are many permutations for RU allocations and the RU distribution across users.
- A mitigation technique influences the SINR per RU, which influences MCS and thus packet airtime.
- The airtime per user should be balanced to avoid excessive padding.

Depending on the mitigation technique used, it can still be applied even if CRUA does not detect any ongoing interference to overcome CTI that starts during a packet transmission, or when the CTI level at the AP was too low to be detected. For example, using CTI feedback from the stations, RU reallocation can be applied without cost when no CTI is present. Power boost factor (i.e., a mitigation technique to boost part of the RU's transmission power) can be applied if the SNR for the unaffected users is already sufficiently high, or when the packet length requirements allow assigning a lower MCS to those users without increasing the packet airtime.

2.3.4 IMPLEMENTATION OF THE FINAL POC AND RESULTS

In this section we present how the techniques for jammer mitigation are implemented on the openwifi platform, as well as the experimental validation and evaluation process followed to evaluate the performance of the PoC.

2.3.4.1 IMPLEMENTATION OF NEW COMPONENTS ON OPENWIFI POC

We use openwifi [13], a full-stack IEEE 802.11 transceiver running on a System-on-Chip, consisting of a Field Programmable Gate Array (FPGA) and ARM processor. The baseband processing and lower MAC-layer is implemented on the FPGA, whereas the driver and higher-level MAC run as Linux kernel modules on the ARM processor. In our previous work [14], we extended it with OFDMA support. When running as AP, the RU allocation can be controlled through defined API of the driver, which allows us to implement the RU reallocation technique.

CRUA is implemented on the FPGA by expanding the use of the existing 256-point FFT core used in the FPGA for decoding the IEEE 802.11ax data symbols. When the receiver is idle, the FFT core is used to calculate the average squared magnitude of the subcarriers within an RU, i.e. its average power. Then, when considering a certain RU allocation, it determines whether one of the RUs has a considerable higher average power than the others.

If the threshold of CRUA is met for one of the RUs, this will be signalled to the transmitter module on the FPGA, such that upon preparing the packet, it can puncture this RU. Puncturing is implemented by denoting the STA identifier to 2046 (as required by the standard) and nullifying the corresponding subcarriers. Furthermore, a scaling factor is applied on the remaining RUs following Equation 1 to keep the total transmit power constant before and after RU puncturing.

2.3.4.2 ADDITIONAL EVALUATION OF RU REALLOCATION

In D3.3, we briefly evaluated the performance of RU reallocation as compared to a naive scheduler, as well as a scheduler that only uses SU packets, which was used to indicate the potential benefit to use interference classifier in an OFDMA scheduler. The naive scheduler uses a fixed RU allocation that assigns an RU to a user where it

experiences narrowband CTI. In those experiments, we chose to use Zigbee sensor nodes as jammers, which are also referred to as the Low Rate Wireless Personal Area Network (LR-WPAN) devices. The duty cycle of the jammer in the simple experiment of D3.3 is fixed, while in this final PoC, we evaluate the performance of RU reallocation when the duty cycles of CTI is varied between 0% to 100%.

Knowing that it is very hard to control CTI duty cycle with real communication, since the duty cycle of the interference may also change depending on Wi-Fi network's activity. This is because Zigbee by default uses Carrier Sense Multiple Access / Collision Avoidance (CSMA/CA). We hence generate the waveform of an IEEE 802.15.4 packet in MATLAB, which is replayed by the R&S CMW270 Wireless Connectivity Tester. The signal is sampled at 20MHz and transmitted using a carrier frequency of 2412 MHz. We therefore shifted the signal to IEEE 802.15.4 channel 11, which overlaps with the first 106-tone RU of Wi-Fi channel 1. The AP uses a fixed MCS 7 to avoid the influence of the rate control algorithm on the measurements.

To create different duty cycles, the repeatedly transmitted IQ samples are appended with a period of zeros. The output power of the tester is set to -60dBm, in order to establish CTI only at STA2, while both STAs are in close proximity to the AP to compensate for its limited transmit power. ESP32-C6 devices are used as Wi-Fi stations, and the openwifi AP has prior knowledge of the RU allocation it should use. This is because under high CTI duty cycle, Wi-Fi PER deteriorates, which affects the CTI feedback towards AP.

One of the stations is placed right next to the tester's output port, while the other is placed on the other side around 0.5m away, with the AP in the middle. We use iPerf to transmit UDP packets using the default payload length of 1470 bytes to the two STAs simultaneously. Five measurements of one minute are performed for each scheduler and the average throughput as reported by the STAs is noted. See Figure 21 for the average throughput for each STA with error bars showing standard deviation, influenced by CTI at different duty cycles. The solid lines show the throughput for the STA that is further away (STA1) from the interferer, while the dotted lines correspond to the STA that is close to it (STA2).

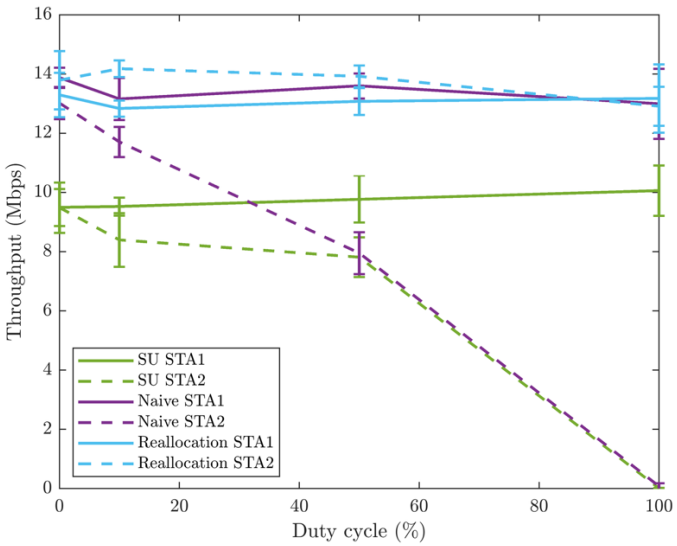

FIGURE 21. DOWNLINK THROUGHPUT MEASUREMENT OF DIFFERENT OFDMA SCHEDULERS FOR DIFFERENT CTI DUTY CYCLE.

First of all, the throughput when using only SU is lower, which can be partly attributed to the increased overhead from the preamble and additional contention time. Furthermore, it can be observed that, compared to the case with no CTI (0% duty cycle), the throughput towards STA2 decreases for both SU and the naive scheduler, meaning these techniques cannot effectively mitigate the CTI. This decrease is more pronounced at higher duty cycles, reaching nearly 0 Mbps at a 100% CTI duty cycle. In contrast, when using RU reallocation, the throughput to both STAs remains roughly constant across different duty cycle settings. This proves that RU reallocation is highly effective for mitigating CTI in scenarios with unequal interference power levels for the STAs.

2.3.4.3 EVALUATION OF RU PUNCTURING

For the influence of RU puncturing on both Wi-Fi and LR-WPAN traffic, we refer to our previous work [15], where we experimentally showed the performance of punctured single-user uplink packets using trigger frames. While in [15] it is assumed that interference is always present, here we demonstrate the capability of CRUA for real-time RU puncturing, namely an RU is only punctured if an ongoing LR-WPAN packet is detected on this RU. We will evaluate the real-time performance of CRUA for RU puncturing in terms of the required detection and scheduling latency.

To assess the real-time performance, it is important to consider the HE PPDU (High-Efficiency Physical layer Protocol Data Unit) structure as used by IEEE 802.11ax in Figure 22.

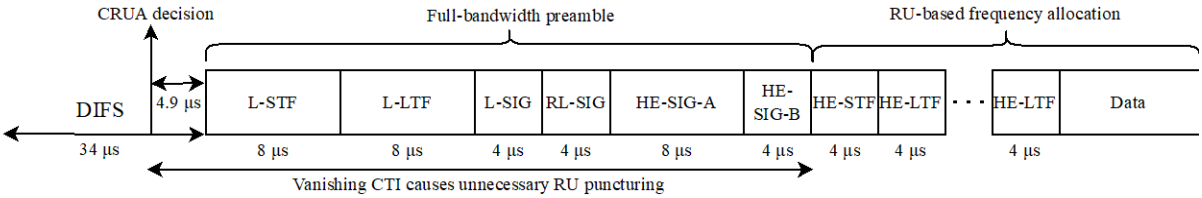

FIGURE 22. WI-FI HE PPDU STRUCTURE WITH DURATIONS PER FIELD AND ANNOTATIONS RELEVANT FOR REAL-TIME PERFORMANCE ASSESSMENT OF CRUA FOR RU PUNCTURING.

In terms of detection latency for the puncturing decision, at a sampling rate of 20MHz, Wi-Fi 6 requires 256 samples for an FFT operation, hence the duration to collect samples for one FFT is 12.8 us. Afterwards, the FFT operation and CRUA calculation take 29 us. Thus, in total when a CTI packet has started at least 41.8 us before preparing a Wi-Fi packet, it can be avoided by RU puncturing. Since the DIFS is 34 us, and the inter-packet gap becomes even longer when random back-off is performed, this usually gives enough time to perform CRUA in between two Wi-Fi packets, even in a saturated channel.

When the CRUA decision is taken, the to-be-transmitted packet has to be prepared and delivered to the RF front-end. The time between the decision and the actual start of the packet on the air is measured to be 4.9 us as denoted in Figure 22. Since the start of the packet still occupies the full bandwidth, RU puncturing has no effect yet there. Using one HE SIG-B symbol, which suffices for the case of two-user multi-user (MU) packet, the full-bandwidth preamble takes 36 us. So, if the CTI packet has already stopped within 40.9 us after the CRUA decision, RU puncturing would be unnecessary, as the CTI does not last till the Wi-Fi data part. Considering that narrowband packets like LR-WPAN can take from around 0.5 ms up to 50 ms, we judge that the chance of unnecessary puncturing is low.

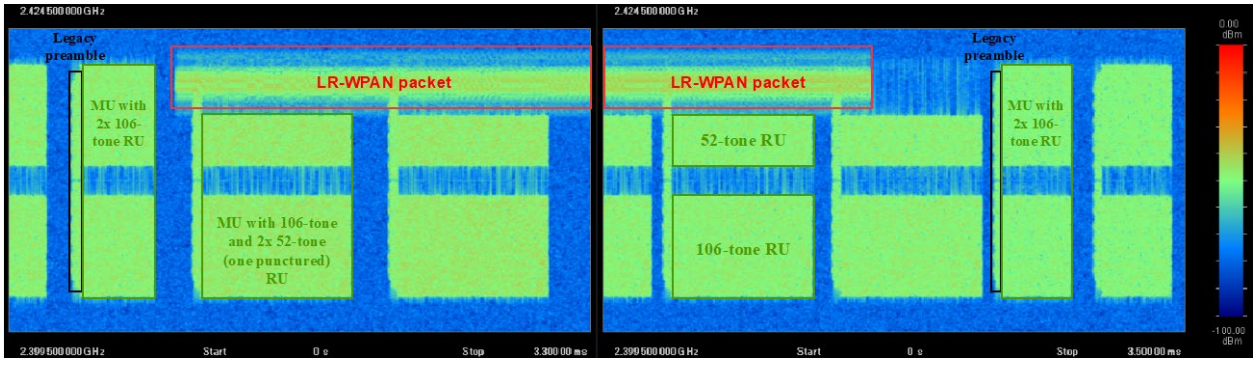

FIGURE 23. SPECTROGRAMS CAPTURED DURING CRUA OF TWO 106-TONE RU ENABLING REAL-TIME RU PUNCTURING.

Next, we use the Anritsu MS2690A spectrum analyser to capture packets transmitted by the openwifi AP to the two STAs, while simultaneously the two LR-WPAN devices are exchanging packets. The openwifi AP operates with 20MHz bandwidth at Wi-Fi channel 1, which centres at 2412 MHz, and the LR-WPAN transmitters use channel 14, which centres at 2420 MHz. The spectrum analyser is set to a centre frequency of 2412 MHz with a span of 25 MHz. Figure 23 shows the captured spectrogram when a LR-WPAN packet starts transmitting (left spectrogram) shortly before the AP starts its downlink transmission. The right spectrogram shows when the LR-WPAN packet

ends. It can be seen that immediately after CRUA detects the CTI, it switches from a 2x 106-tone MU packet to one with 1x 106-tone and 2x 52-tone RUs, among which one 52-tone RU is punctured to avoid LR-WPAN interference. Then, when the LR-WPAN packet ends, the scheduler switches back to the original RU allocation for the next packet. This shows that the real-time RU puncture is successful.

It can be seen that the legacy preamble is still transmitted on the full 20 MHz bandwidth for backward compatibility. Thus, this still causes a short period of interference on the LR-WPAN packet. The work in [11] showed the impact of the Wi-Fi preamble on ZigBee transmissions. First, they argue that when it interferes with the ZigBee preamble, which is 128 us long, it does not affect the ZigBee transmission much due to the redundancy of its preamble. When the Wi-Fi preamble interferes with a ZigBee data symbol long enough, it will affect its decoding ability, but [11] shows this can be effectively mitigated by applying Reed-Solomon error correction codes in the ZigBee packets. Our previous work [15] also show that even with RU puncturing, Wi-Fi still has an impact on LR-WPAN PER, due to the full bandwidth preamble. On the other hand, the Wi-Fi legacy preamble is thus also interfered by LR-WPAN, but since the fields in the Wi-Fi preamble use the most robust modulation and coding scheme, and channel estimation for the data part is performed using the partial-bandwidth HE preamble fields, the impact of Zigbee on Wi-Fi is acceptable.

2.3.4.4 CONCLUSION

In this PoC, we use Wi-Fi6 compliant OFDMA scheduling to improve physical layer robustness against narrow band jammer. In particular, two techniques are evaluated, i.e. RU reallocation and RU puncturing. RU reallocation brings significant benefits when there is a large interference level difference experienced by the Wi-Fi stations, e.g., almost no throughput loss is observed when interfered compared to the 50% throughput drop when the OFDMA based jammer avoidance is not active. RU puncture on the other hand can be applied when all Wi-Fi stations experience similar level of interference, then instead of full bandwidth, partial bandwidth can be applied when Clear Resource Unit Assessment (CURA) identifies the jammer.

The techniques (RU puncture and reallocation) are standard compliant however, to take effective decision CURA and CSI based interference detection and feedback mechanisms are also necessary. These are the key innovations to support the robot control use case with Wi-Fi supported sub networks of 6G technology. The PoC also shows that the jammer classification based on channel state information in D3.3 can be effectively applied in real-life.

2.4 INTRA-SUBNETWORK MACRO-DIVERSITY

2.4.1 GENERAL DESCRIPTION

In 6G-SHINE, FHG contributes to TC8: Intra-subnetwork macro-diversity. In this context, methods to exploit spatial diversity via the PHY multi-links and via subnetwork devices in range by using cooperative mechanisms based on network coded cooperation (NCC) are investigated.

Such macro-diversity solutions can be used in industrial subnetworks for achieving robustness to blockage effects that may hinder the efficient support of critical services. It is a very likely scenario that installations in production cells, machines or robots, like the unit test cell use case identified in D2.2 [1], may lead to situations where the desired link between a base station (BS) which is a HC device and multiple SNEs such as sensors, actuators or robots is obstructed by obstacles such as metallic items. Since programmable logic controllers (PLC) typically work in deterministic cycle times, not only the latency of a transmission but also the round-trip time of a communication becomes very important.

The challenge therefore is to design a transmission and coding scheme optimized in terms of challenging KPIs – especially the packet loss rate - for subnetworks with multiple devices. Innovative Multi-Link Operation (MLO) based on NCC concepts were developed and evaluated by simulation in WP3 (see Deliverable D3.3 [1]). The simulation results show outstanding improvements of a network reliability, especially in the case of obstacles in the line of sight between HC and SNE.

WP2 analysed the requirements for industrial subnetworks use cases and reported them in Deliverable D2.2. For our PoC - like for our simulation work in WP3 - the relevant requirements were implemented as follows:

Requirements from D2.2	**Configuration of the communication parameters used in PoC**
Cycle time < 100 µs..1 ms for communication in loops other communication cycles 10..100 ms	Cycle time is 250 µs, which corresponds to a TDMA frame duration of 250 µs, comprising 26 time slots. We assume symmetric bidirectional communication between HC and SNE. Data packets should be delivered within a single frame (deterministic cycle time!)
Data packet size 100..300 Byte	28 Byte of higher layer data The data packet size could be extended by segmentation and using multiple resource blocks, but it was considered as being sufficient to evaluate the error control mechanisms.
Data rate 1..2 Mbit/s (in case of video up to 80 Mbps)	With the selected data packet size, the resulting data rate for higher layer data is 896 kbit/s.
20..50 SNE	Predefined and stationary set of one HC and 2 SNE
Movement speed of nodes none or up to 20 m/s	Stationary nodes but different topologies and varying obstacles
No requirement defined	Multilink operation: two parallel links between devices (operation on second channel synchronous to first channel (same slot and frame alignment)) This operation mode allows shorter transmission phases und thus better NCC strategies.
No requirement defined	Star topology: no connection on application level between SNE, only between HC and SNE
probability of having two consecutive errors < 10^{-6}	This will be the outcome of the measurements.

TABLE 5. COMPARISON OF THE REQUIREMENTS WITH THE IMPLEMENTATION IN 6G-SHINE

The objective of our PoC system is the verification of the simulation results by an over the air communication performed in real time based on a prototype implementation.

2.4.2 ENHANCED/NEW FUNCTIONALITY VS PREVIOUS VERSION

The PoC system realizes a TDMA/TDD communication with an ultra-low cycle time of 250 µs. HC and SNEs support PHY multi-links which allow simultaneous transmissions in two frequency bands. Each SNE has the ability to overhear packets which are originating from or are direct to other SNEs, which is one important prerequisite for cooperative communication. This is crucial since it enables the PoC system to benefit from path diversity.

Due to unforeseen challenges in the real-time implementation, a decision was made to only implement a subset of the promising Network Coded Cooperation (NCC) strategies reported in Deliverable D3.3. The algorithms for encoding and decoding were fully implemented, which provides each node with the ability to build coded packets from native packets and to recover native packets from coded packets. This functionality is used for the implementation of the NCC strategy in which each SNE is supporting exactly one other SNE by sending out a coded packet which combines the Uplink and Downlink packet for the SNE it wants to support. The assignment between supporting SNEs is static by the usage of static coding vectors.

2.4.3 HIGH LEVEL ARCHITECTURE CHANGES/UPDATES

The functionality of the PoC system is achieved through the interaction of numerous hardware and software components. In terms of the architecture (see Figure 24) there is no difference between a SNE and the HC.

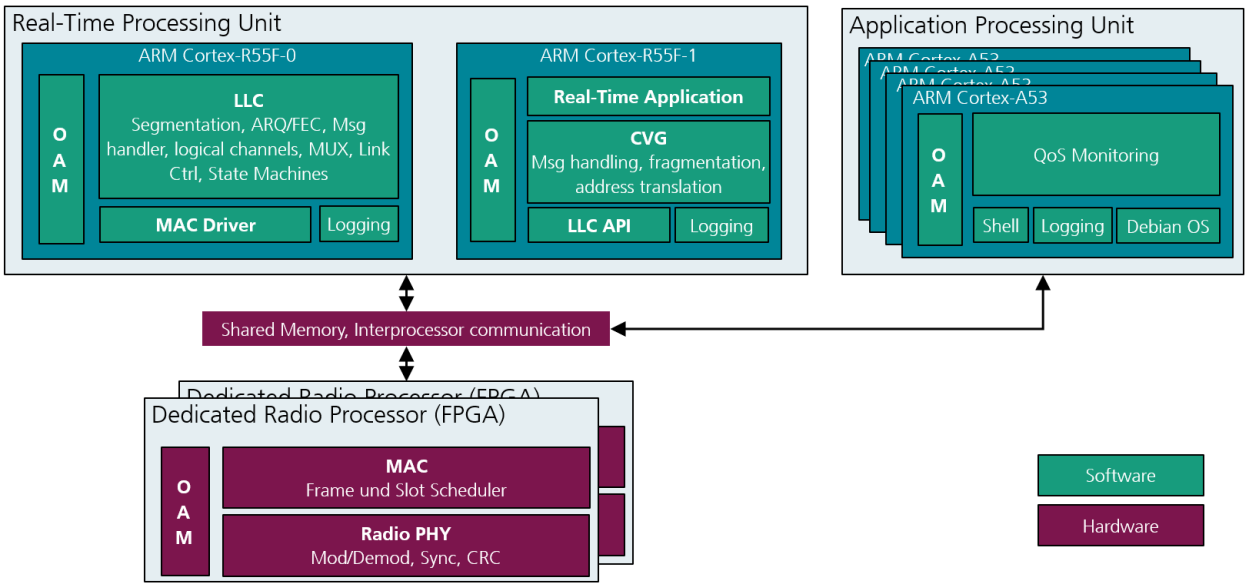

Figure 24. Architectural diagram for Software and Hardware Components of HC and SNE

The Application Processing Unit (APU) is used for a defined startup and configuration (Operation and Management, OAM) of the radio node. The QoS Monitoring Module allows to collect and analyse key performance indicators like cycle times and packet loss rates during operation.

The Real-Time Processing Unit (RPU) hosts a Real-Time Application for processing and cyclic exchange of process data. The Convergence (CVG) Layer is an optional layer which is used for protocol translation and the fragmentation of large packets into a fragment size which the Logical Link Control (LLC) Layer supports. The LLC is multiplexing/demultiplexing logical channels and applies error control methods (see previous subsection) on process data. It also derives QoS metrics and used dedicated logging channels to share this information with the APU.

For the purpose of multi-link communication, two instances of the PHY/MAC layer are available.

For low-latency interprocessor communication, the hardware and software components require a dedicated shared memory. The shared memory is implemented as a Block RAM in the FPGA which allows the RPU an access with access times below 100 ns.

2.4.4 FINAL DEMONSTRATOR SCENARIO DESCRIPTION UPDATES/CHANGES

The demonstrator scenario (see Figure 25 and Figure 26) was selected to investigate the effect of obstacles and their impact on the reliability of the communication in subnetworks in particular with regards to industrial use cases. For this purpose, a moving obstacle and a network of stationary radio nodes are considered. In the basic scenario two SNE are in close proximity to one HC (Figure 25). The HC and each SNE are attempting to exchange 28 Bytes of application data with a cycle time of 250 µs.

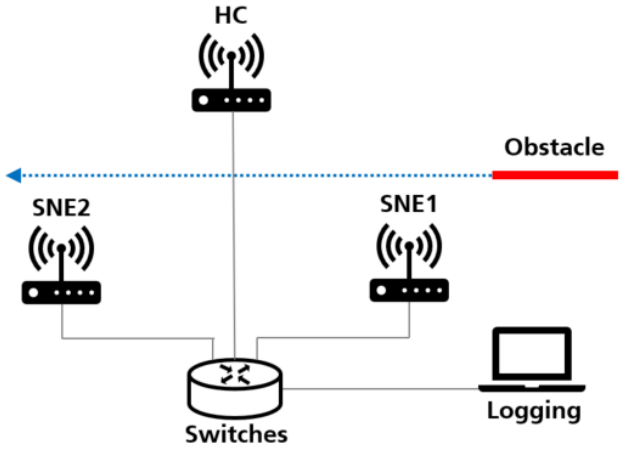

Figure 25. Basic Demonstrator Scenario

To prove that packet loss rates below 10^{-6} can be achieved, 10^{8} cycles are recorded in the following three spatial setups:

- Setup 1: Obstacle is exclusively impairing the line-of-sight connection between HC and SNE 1
- Setup 2: Obstacle is not impairing any line-of-sight connection
- Setup 3: Obstacle is exclusively impairing the line-of-sight connection between HC and SNE 2

Figure 26 gives an impression on the setup's realization, as it is complicated to move an attenuation obstacle into the line-of-sight condition the disturbance is achieved by moving SNEs further away behind concrete walls.

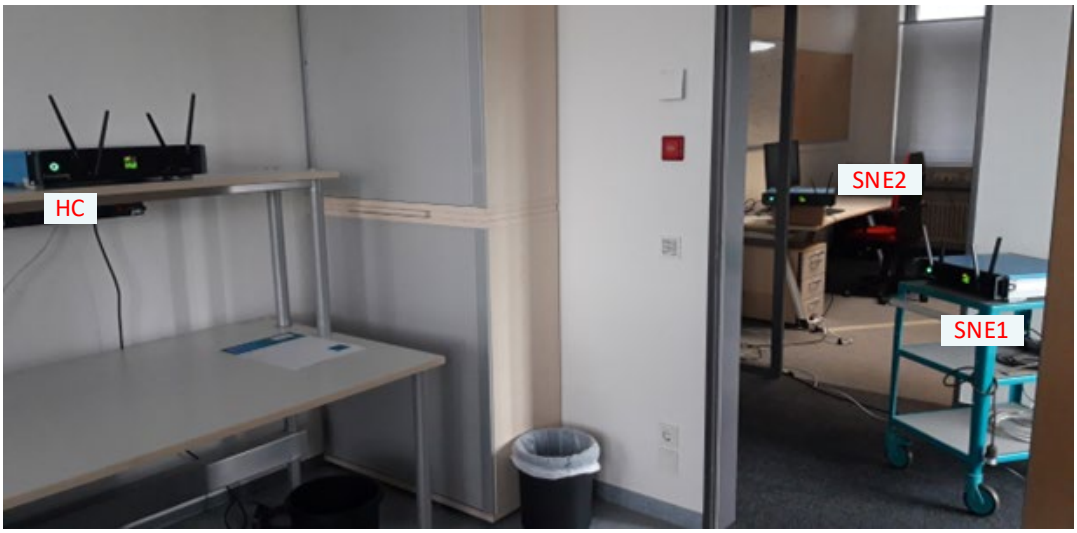

Figure 26. Demonstration Setup realization

All radio nodes are connected via a switch to a dedicated PC which collects the derived QoS metrics of each node in the network.

2.4.5 RESULTS

2.4.5.1 UPDATED TC RESULTS VS THE BASELINE AND/OR VS VERSION 1.0

Debugging has proven to be a challenging task, primarily because observing the system can influence its behaviour, and the presence of multiple and distributed timebases complicates the process further. Additionally, different tools are required for debugging software and hardware, which adds to the complexity. Resource management

poses its own difficulties; the cores designated for real-time processing have very limited memory available for program code, necessitating careful selection and configuration of used libraries. The complexity of the platform architecture adds another layer of challenge. Furthermore, the requirement for ultra-low cycle times has led to the decision not to use an operating system on the real-time processing unit, which increases the effort needed for resource management, task scheduling and logging.

These challenges are typical for real-time applications that are partially implemented in hardware and software, and they were anticipated, but the effort required to address them was not known by the time the project proposal was prepared. For this reason, not all packet loss rates that were initially planned to be measured could be obtained (see Table 6).

Table 6. Measured packet loss rates (PLR) for the NCC error control method

Spatial setup		1		2		3	
Resource allocation strategy		Single Transmission	NCC	Single Transmission	NCC	Single Transmission	NCC
PLR SNE1	UL	-	-	-	-	-	-
	DL	-	-	$2.0 \cdot 10^{-9}$	-	-	-
PLR SNE2	UL	-	-	-	-	-	-
	DL	-	-	-	-	-	-

The application of NCC does not have an influence on the cycle time of 250 µs. Since the packet error rates under line-of-sight conditions were lower than expected, we conducted a measurement over $5 \cdot 10^{8}$ cycles, resulting in a single packet loss (leading to a $2.0 \cdot 10^{-9}$ loss rate).

2.4.5.2 WP2 USE CASE MAPPING OF ACHIEVED RESULTS AND WP2-3-4 THEORETICAL RESULTS

We can realize communication with a very low cycle time of 250 µs and achieve this important requirement. The communication under line-of-sight conditions is also ultra reliable and therefore we expect that we can also achieve the required reliability under non line of sight conditions using NCC. With a maximum achievable data rate of 896 kbit/s, the required data rate of at least 1 Mbit/s was almost achieved. These results match the simulation results in D3.3. The PoC system currently supports only one SNE which is not sufficient for most industrial use cases. However, since the radio resource usage in the PoC is very little and the cycle time is only 250 µs, there is still a lot of potential to scale up the supported SNE amount.

2.4.5.3 CONCLUSION

In 6G-SHINE WP3 macro diversity methods for exploiting spatial diversity via the PHY multi-links and via subnetwork devices in range by using cooperative mechanisms based on network coded cooperation (NCC) have been developed. Such macro-diversity solutions can be used in industrial subnetworks for achieving robustness to blockage effects that may hinder the efficient support of critical services. It is a very likely scenario that installations in production cells, machines or robots may lead to situations where the desired link between a base station (BS) which is a HC device and multiple SNEs such as sensors, actuators or robots is obstructed by obstacles such as metallic items. In simulation Network Coded Cooperation shows excellent results in 6G subnetworks to increase the availability of SNEs and the reliability of the transmission. NCC can also be used in TDMA/TDD systems with extremely short cycle times. In this project opportunistic Network Coded Cooperation was investigated, which is a new method where an evaluation of the connection quality is carried out in the SNE to identify which other SNE can / should be helped in the NCC phase.

A PoC was set up to test the processes in real time and in a real environment.

A very powerful prototype hardware based on FPGA, real-time processing units and application processing units was set up and TDMA/TDD-based protocol processing for PHY, MAC and LLC was realized.

The cycle time of the PoC is 250 µs for bidirectional transmission including error control. This makes the PoC approximately 10 times faster than today's 5G URLLC systems.

The PoC was used to carry out initial transmission tests with logging of packet error rates for single transmission. In line-of-sight conditions a packet error rate of $2.0 \cdot 10^{-9}$ was measured, which impressively demonstrates the possible transmission performance of the implemented PoC.

The complexity of the real-time processing meant that the implementation of the error control procedures is still ongoing. The measurements to determine the packet loss rates for NCC could not be completed by the time this report was submitted. They are scheduled to be carried out by the end of the project period and the results will be published afterwards. Even with this implication, we were still able to draw very important conclusions for the feasibility and possible impact of the relevant TC and we are now sure that the impact of the proposed TC will be as significant as expected based on the theoretical analysis initially but now also based on the results of this PoC.

2.5 WIRELESS IN-VEHICLE E/E ARCHITECTURE

2.5.1 GENERAL DESCRIPTION

In Deliverable D2.4 of WP2 [1], the Non-3GPP Subnetwork Management Function (N3SMF) was initially introduced as a key architectural element to support integration between subnetworks and the parent network. Its definition included the functional role, interfaces, and associated protocols, forming the foundation for several use cases outlined in D2.2 [1], most notably V-3 (Inter-subnetwork Coordination for Intravehicular Communication) and I-4 (Subnetwork Co-existence in Factory Hall), as both use-cases directly require inter-SN coexistence and resource management among different entities.

In this deliverable, we present the full implementation of the N3SMF component. This work builds upon the concepts and initial developments presented in D5.2 [1], extending them into a fully operational and integrated solution. More specifically, the N3SMF was successfully integrated into the parent network infrastructure, enabling real-time subnetwork monitoring and execution of RRM procedures. Furthermore, the implementation utilizes contextual data, such as high-accuracy ML-based positioning, to proactively mitigate inter-subnetwork interference scenarios.

In more details, we describe the full implementation of the proposed framework, provide results from our PoC, and discuss future directions. The PoC is structured around three key stages, each of which is detailed below:

1. We establish a wireless in-vehicle communication infrastructure by adapting elements of the E/E architecture to support wireless transmission. This results in a dedicated subnetwork managed by a Subnetwork Controller (SNC).
2. The in-vehicle subnetwork, based on short-range, low-power wireless connectivity, is integrated into a broader parent network to enable coordinated resource sharing and management.

3. Within the parent network, we implement mechanisms to dynamically allocate communication resources across multiple in-vehicle subnetworks. This is achieved by applying contextual data, such as real-time vehicle positioning, to minimize interference and optimize performance.

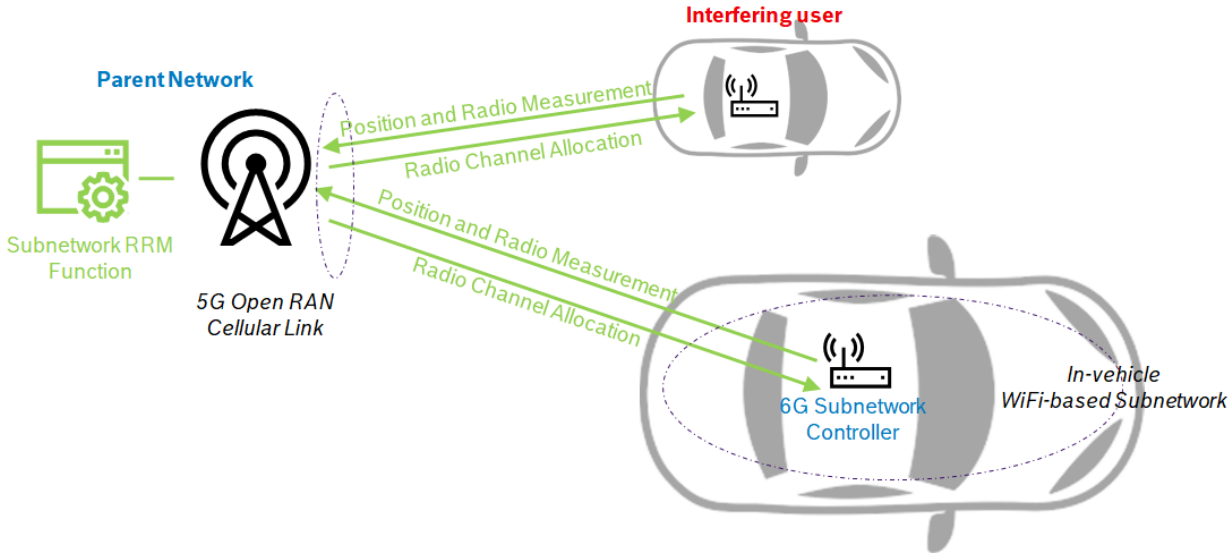

FIGURE 27. OVERALL POC DESCRIPTION

The overall scenario, illustrated in Figure 27, depicts a 6G-enabled parent network that orchestrates interference management across several in-vehicle subnetworks. For the in-vehicle subnetwork implementation, we adapted the automotive CAN protocol for Wi-Fi using Commercial Off-The-Shelf (COTS) hardware running OpenWRT [21]. In this configuration, it is important to note that latency performance may be influenced by the listen-before-talk (LBT) mechanism [22]. To achieve optimal performance, vehicles in close proximity should ideally operate on orthogonal frequency resources, thereby minimizing mutual interference.

2.5.2 ENHANCED/NEW FUNCTIONALITY VS PREVIOUS VERSION

In Deliverable D5.2 [1], we presented results using automotive radar to demonstrate the effectiveness of our proposed positioning-based Radio Resource Management approach. Building on this, we now extend the concept to a different but related application: wireless automotive cameras. Specifically, we focus on their use in surround view systems powered by stitching algorithms [28], as illustrated in Figure 28. This application presents a compelling use case for three main reasons:

1. It better reflects the strict network traffic requirements outlined in WP2 D2.2 [1] (refer to Table 27) particularly for Advanced Driver Assistance Systems (ADAS) related applications.
2. In certain countries, most notably China, surround view systems are being mandated by national regulations [23].
3. 6G in-X subnetworks can serve as a competitive alternative to emerging technologies such as NearLink, which is currently gaining traction in China [24].

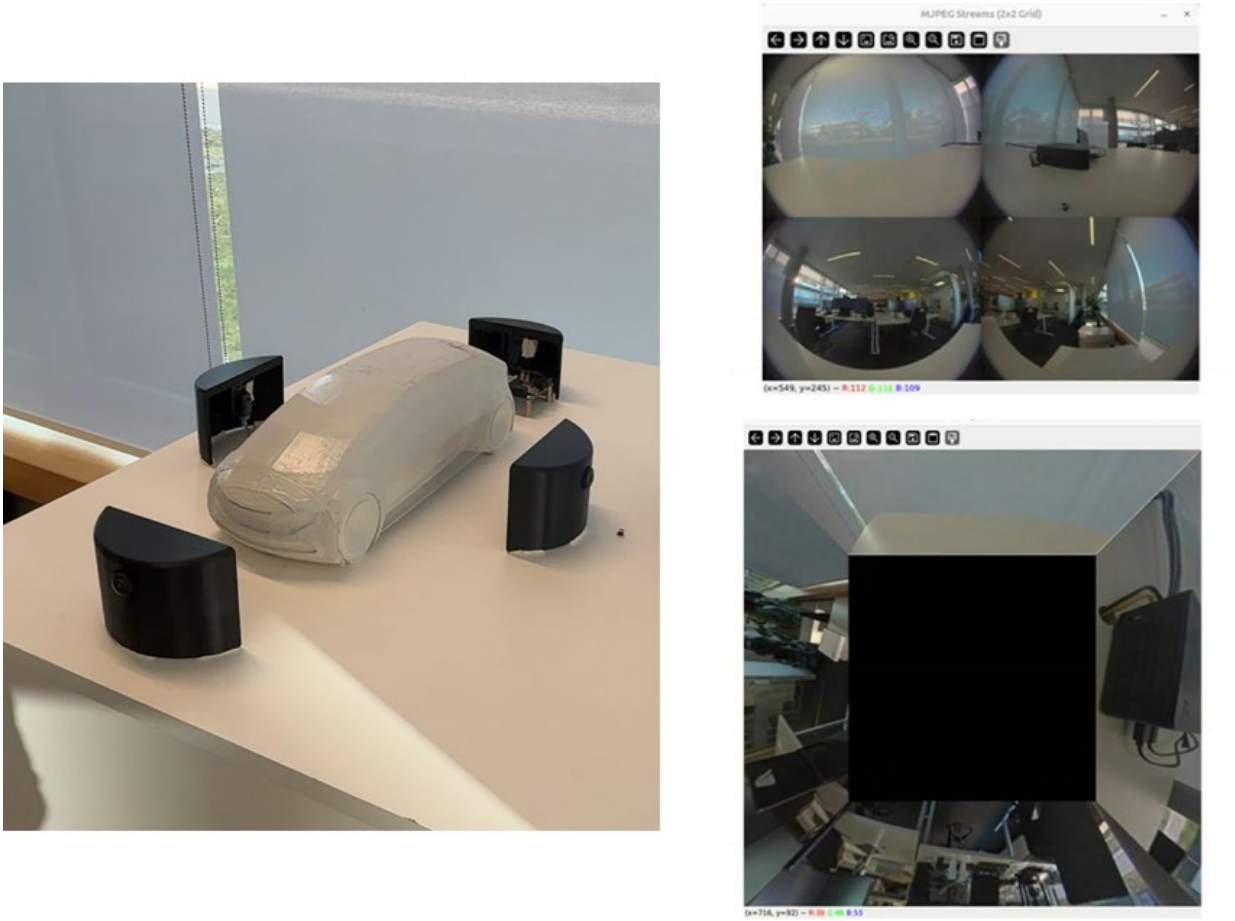

FIGURE 28. SURROUND VIEW CAMERA SYSTEM USING STITCHING ALGORITHM.

As with D5.2, we establish a parent network that assumes a central RRM role, including monitoring of non-3GPP networks. This is realized using the proposed Parent Network (PN) feature known as the Non-3GPP Subnetwork Management Function (N3SMF). The N3SMF enables active monitoring and management of Wi-Fi-based services through the 5G network by exposing relevant functions to the 3GPP architecture. Consequently, the 5G Core Network (CN), based on the Service-Based Architecture (SBA), can perform centralized RRM tasks across heterogeneous networks.

We also incorporate an indoor positioning solution based on Ultra-Wideband (UWB) technology, utilizing a Graph Neural Network (GNN) to infer positions with sub-5 cm accuracy, even without line-of-sight to all anchors. This method has been accepted for publication at IEEE Globecom 2024, with further technical details provided in [25]. The computed position data is reported to the N3SMF within the parent network for context-aware RRM.

The overall system architecture is illustrated in Figure 29. The main building blocks follow the standard 5G Service Based Architecture (SBA), comprising the Access and Mobility Management Function (AMF), Session Management

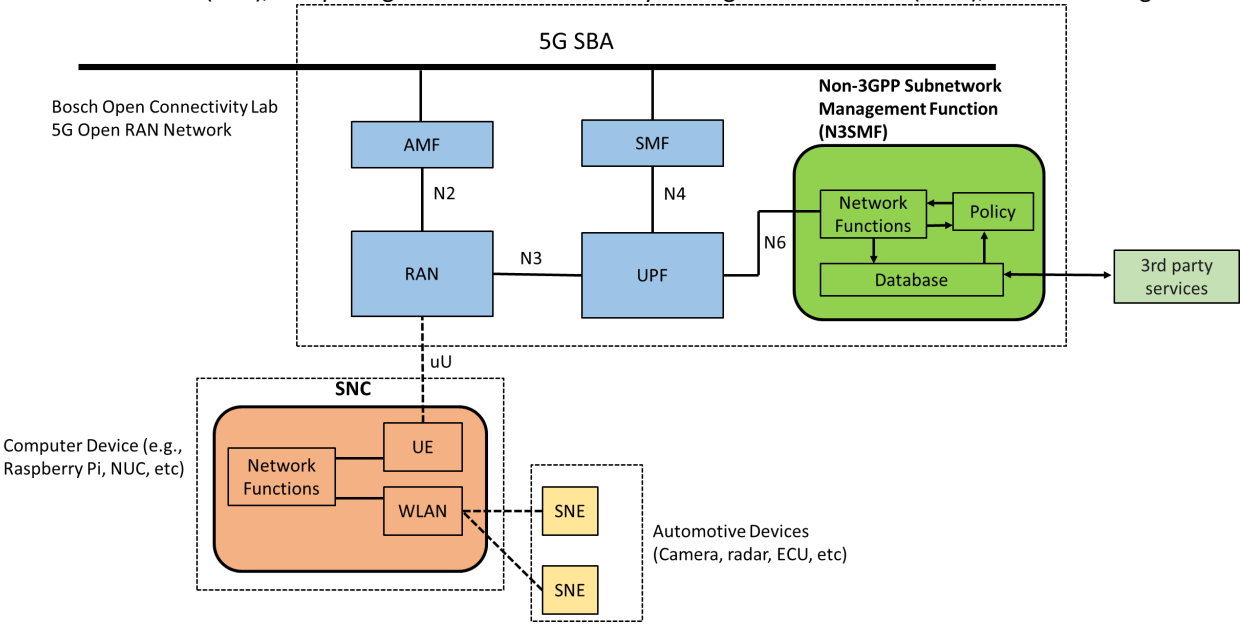

FIGURE 29. INTRODUCTION TO N3SMF AS A KEY ENTITY TO MANAGE NON-3GPP SUBNETWORKS FROM THE CORE NETWORK

Function (SMF), User Plane Function (UPF) and the Radio Access Network (RAN). The N3SMF is developed as an extension on top of the UPF and maintains a virtual interface with the SNC to support RRM decision-making process.

2.5.3 HIGH LEVEL ARCHITECTURE CHANGES/UPDATES

Details of the PoC setup are depicted in Figure 30. In this configuration, we integrate Bosch Near-Range Automotive Cameras [26] as subnetwork elements. These cameras are originally designed to operate using the IEEE 1722 standard; we have modified the protocol to support wireless transmission of camera data.

A Subnetwork Controller manages the wireless communication between the SNEs and the vehicle's central computing unit, where sensor data is processed. In scenarios where other vehicles are nearby, wireless interference can degrade the quality of communication from the SNEs. This, in turn, can impact the integrity and timeliness of the stitching image processing at the vehicle computer. To address this, the SNC receives RRM commands from the parent network, which are triggered by the non-3GPP Sensing Management Function. Automotive camera systems, especially surround view applications, introduce unique challenges for wireless communication due to their stringent latency and time synchronization requirements, particularly in closed-loop control scenarios relevant to ADAS. These characteristics make them a compelling domain for research into wireless performance and optimization. Moreover, the methodology demonstrated here is extensible to other use cases, such as the radar-based scenario explored in the previous deliverable D5.2 [1].

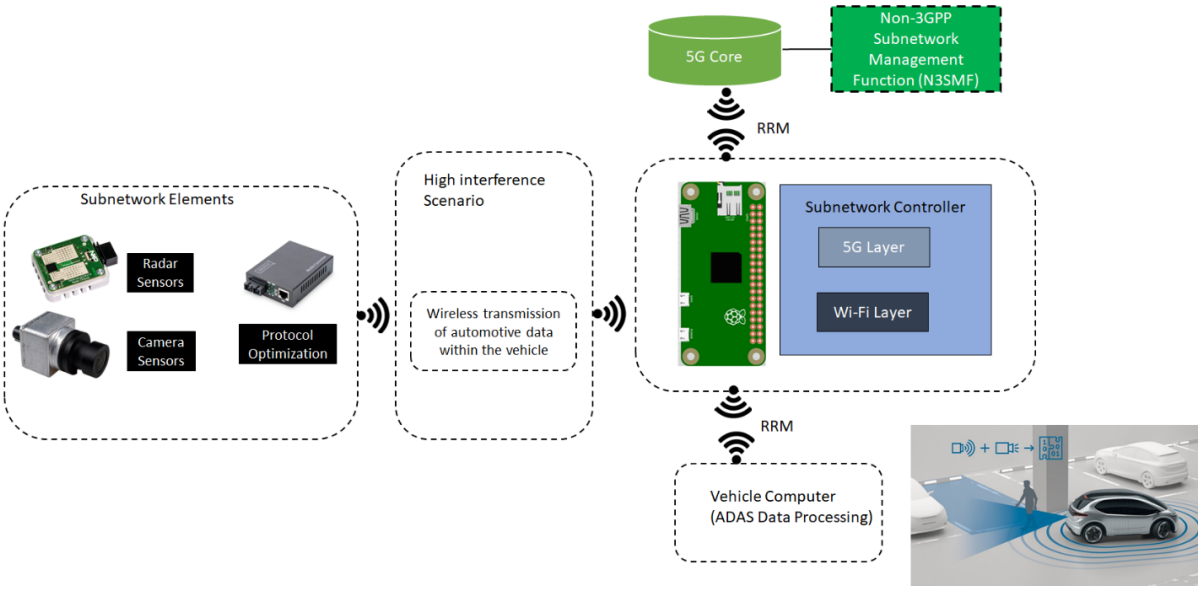

FIGURE 30. HIGH LEVEL ARCHITECTURE OF THE PROPOSED WIRELESS E/E-A POC.

2.5.4 FINAL DEMONSTRATOR SCENARIO DESCRIPTION UPDATES/CHANGES

The scenario demonstration, illustrated in Figure 31, involves configuring automotive components on two adjacent moving tables, each emulating a wireless in-vehicle application, as initially introduced in Figure 29. The goal is to demonstrate the parent network's Radio Resource Management functionality through a two-step process:

1. Initial Scenario (RRM Disabled):

In the first step, the two tables are gradually moved closer together. Their relative positions are tracked by the parent network. At this stage, with RRM disabled, two key outcomes are expected:

- At the parent network level, each SNC will report both the increasing proximity and elevated communication latency.
- On each table, the application displays (e.g., video or radar outputs) will visibly reflect the negative effects of interference -such as lost frames or degraded signal quality.

2. Scenario with RRM Enabled:

In the second step, the same setup is repeated, but this time with RRM enabled. As the tables approach each other, the system detects their proximity and triggers an RRM command. In this case, no performance degradation is expected, as interference is mitigated through dynamic resource allocation. The demonstration also highlights the system's response to the RRM decision from both the network and application perspectives, including whether any data loss occurs during the transition.

The actual implementation of this scenario is shown in Figure 32, where the PoC was internally validated and tested.

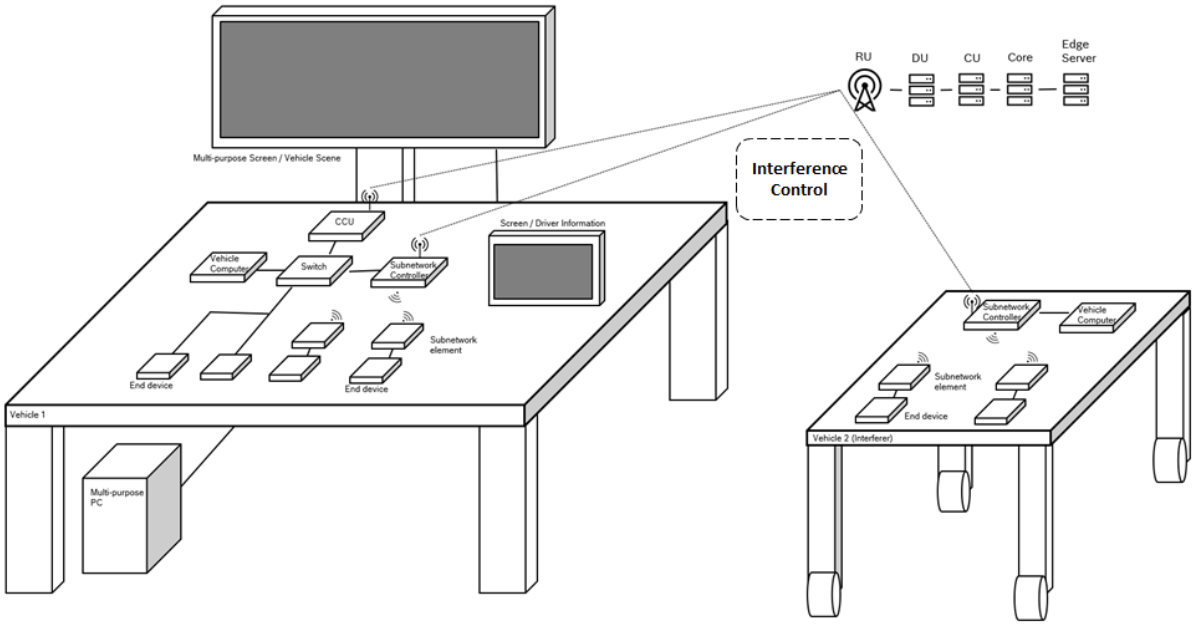

FIGURE 31. OVERVIEW OF THE POC SCENARIO DEMONSTRATION.

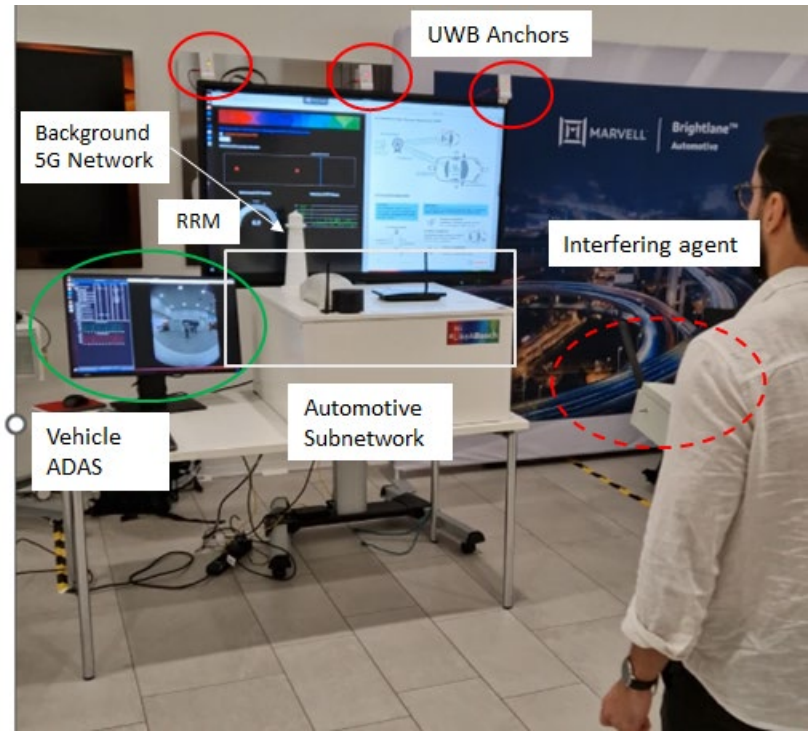

FIGURE 32. DEMONSTRATION IN LAB ENVIRONMENT OF THE PROPOSED WIRELESS IN-VEHICLE E/E -A.

2.5.5 RESULTS

2.5.5.1 UPDATED TC RESULTS VS THE BASELINE AND/OR VS VERSION 1.0

In Deliverable D5.2 [1], we demonstrated the transmission of automotive radar data over a wireless Wi-Fi channel. In that scenario, Wi-Fi performance metrics were monitored in real-time by a 5G Open RAN (ORAN) parent network, which made RRM decisions. These decisions were then communicated to the Subnetwork Controller to mitigate inter-subnetwork (inter-SN) interference.

In this deliverable, we build upon that methodology and introduce two key enhancements:

1. We focus on the wireless surround view system use case, which imposes more stringent network requirements. These requirements are closely aligned with those outlined in WP2.
2. We integrate an indoor positioning system into the parent network infrastructure to support RRM decision-making. The system enhances spatial awareness at the subnetwork level, enabling proactive interference management.

Based on these two enhancements, we compare system performance with and without the parent network's RRM functionality enabled. Specifically, we assess the impact on wireless communication performance within the subnetwork.

To quantify this, we analyse the **inter-arrival time** of packets received at the vehicle's central computing unit. This metric is particularly relevant because each of the four cameras in the surround view system transmits a burst of 5,000 UDP packets per second. Each packet is approximately **8K bits in size**, resulting in a total throughput of around 40 Mbps per camera. Figure 33 illustrates a snapshot of the sniffed traffic captured at the camera output.

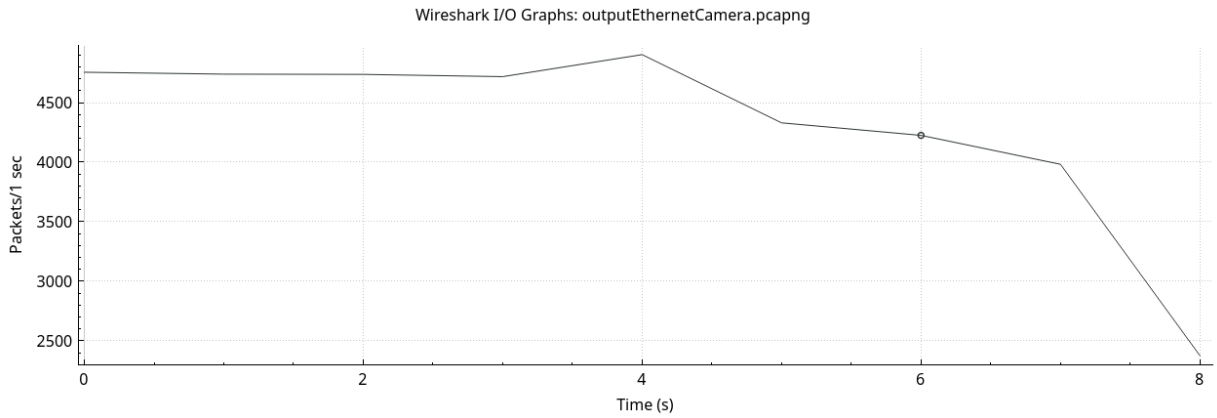

FIGURE 33. CAMERA DATA TRAFFIC WIRELESS TRANSMITTED WITHIN THE E/E-A.

Figure 34 and Figure 35 present the results of packet inter-arrival time measurements in two scenarios: with and without inter-SN RRM enabled. A key feature of our proposed approach is its proactive response to interference. Rather than merely reacting to interference after it occurs, the system anticipates potential conflicts by leveraging contextual information, particularly from the integrated indoor positioning system. This enables the parent network to take preventive RRM actions before interference degrades performance. The PN has visibility into the positions of potential subnetworks and real-time wireless metrics, such as the current channel in use. This contextual awareness allows the RRM function to make intelligent resource allocation decisions while avoiding unnecessary overprovisioning.

In this context, Figure 34 displays the sequence of packet inter-arrival intervals. When RRM is disabled, noticeable peaks in the inter-arrival time appear, indicative of interference that degrades the performance of the stitching algorithm. For instance, with a camera configured for 60 frames per second (FPS), a packet delay exceeding 16 ms between packets can result in a dropped frame, disrupting the machine learning-based stitching process. Looking ahead to future ADAS applications, which may operate at 120 FPS or higher, the system will need to meet even stricter timing requirements. This behaviour is further illustrated in Figure 35, which shows the Cumulative Distribution Function (CDF) of packet inter-arrival times. The results clearly demonstrate that without RRM, a higher proportion of packets exceed critical delay thresholds, leading to more frame losses and degraded application performance.

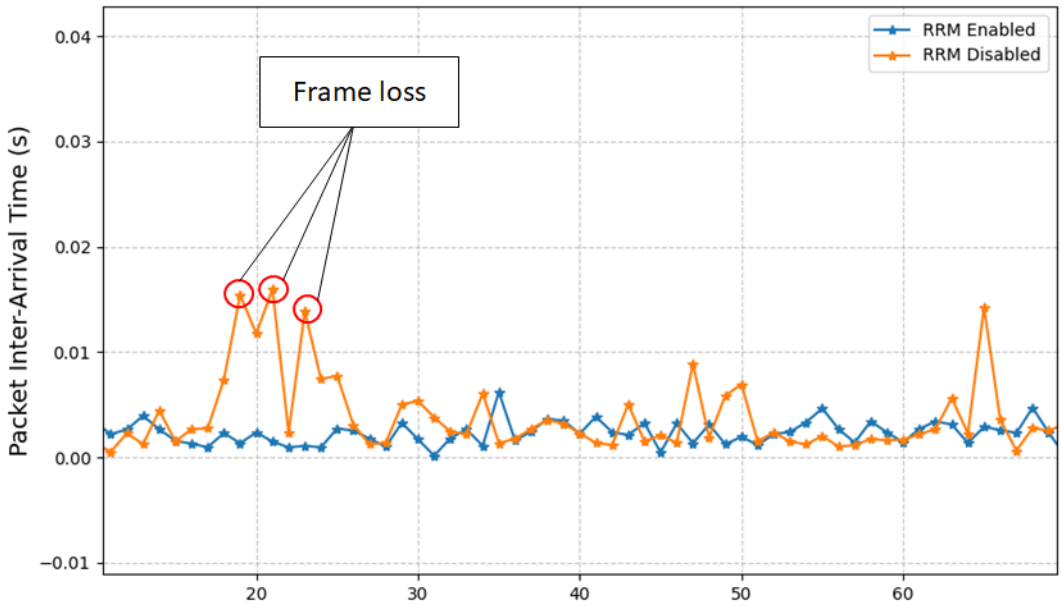

FIGURE 34. SEQUENCE OF PACKET INTER-ARRIVAL TIME RECEIVED AT THE VEHICLE COMPUTER FOR ADAS PROCESSING.

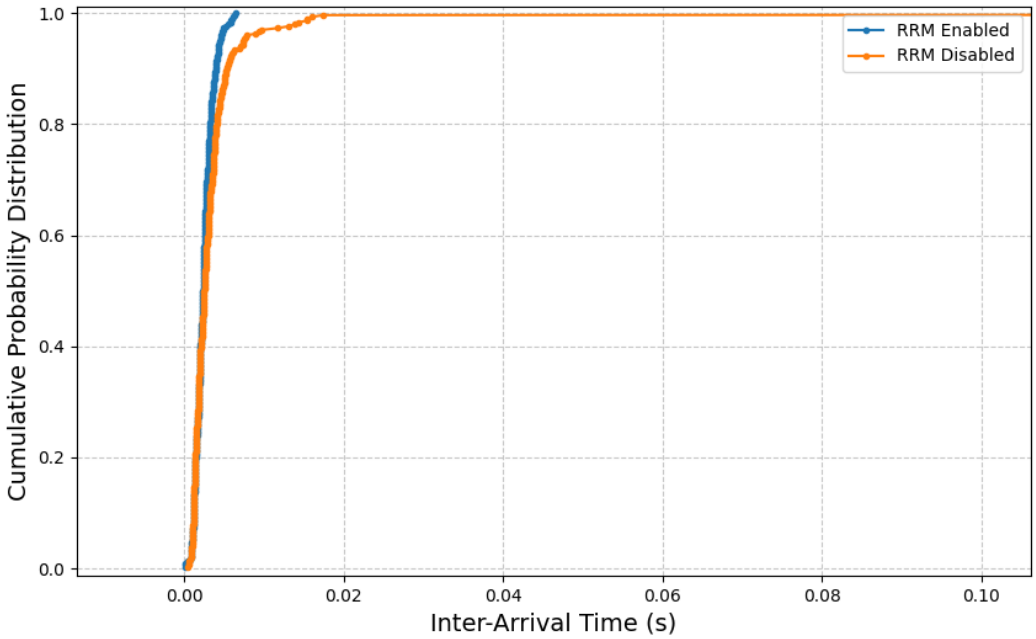

FIGURE 35. CDF OF THE PACKET INTER-ARRIVAL TIME RECEIVED AT THE VEHICLE COMPUTER.

2.5.5.2 ACHIEVED RESULTS VS WP2-3-4 THEORETICAL/SIMULATION RESULTS

We validated the proposed architecture using two representative ADAS-related use cases:

1.A wireless-connected radar system, as detailed in D5.2.

2.A wireless surround view camera system, presented in this deliverable.

Both use cases were evaluated against the communication requirements defined in D2.2, Table 27. The results demonstrate that performance levels, originally specified for traditional wired automotive technologies (e.g., Ethernet, CAN) can also be achieved over wireless networks. This is made possible through RRM strategies, such as the one implemented through the N3SMF framework.

More specifically, our primary objective was to achieve wireless communication performance comparable to that of wired systems, as defined in Table 27 of D2.2, which outlines stringent automotive requirements, including those relevant to ADAS applications.

The system was able to meet the required throughput for the ADAS use case, delivering approximately 120 Mbps across the full surround-view camera setup. Furthermore, we observed an average packet inter-arrival time of around 1 ms, which is within acceptable bounds for real-time ADAS applications.

However, we believe this latency is not primarily the result of radio resource usage or interference but rather stems from protocol-level inefficiencies. Specifically, internal system procedures required to translate and adapt protocol formats appear to introduce overhead that increases packet delay. This observation suggests that protocol and system-level optimizations play an important role in achieving ultra-low-latency performance. Similar behaviour has already been identified and discussed by the authors in [27].

Despite these limitations, the RRM strategies proposed and implemented in this work demonstrate that it is indeed possible and feasible to achieve performance levels comparable to wired systems. This finding reinforces the feasibility of wireless in-vehicle networks for high-demand applications, provided that both radio and protocol layers are co-optimized.

2.5.5.3 CONCLUSION

In this PoC, we presented the full implementation of the Non-3GPP Subnetwork Management Function, initially defined in D2.4 [1]. This implementation enabled the successful integration of in-vehicle subnetworks into a parent 5G network infrastructure and established a method for real-time monitoring and management of subnetwork radio resources. Critically, the solution incorporates contextual information, such as high-accuracy positioning, to proactively address inter-subnetwork interference scenarios, rather than reacting to interference after it occurs.

The proposed system was validated using two automotive applications aligned with the use cases defined in D2.2: First, a wireless-connected radar system, as demonstrated in Deliverable D5.2. Second, a wireless surround view camera system, as explored in this deliverable.

Both applications are used in the context of ADAS and were evaluated under the communication requirements specified in D2.2. The results confirm that these strict performance requirements, originally intended for wired technologies such as Ethernet and CAN, can also be met using wireless solutions, provided that advanced RRM mechanisms like the ones proposed here are in place. Our key conclusion is that wireless-connected automotive components can achieve performance parity with their wired counterparts when supported by intelligent RRM strategies and context-aware network management. As demonstrated, the application of the N3SMF and the proactive RRM framework allows us to meet 100% of the critical requirements outlined in D2.2 for ADAS related applications, establishing a foundation for future 6G in-vehicle subnetwork deployments.

2.6 CENTRALIZED/DISTRIBUTED INTERFERENCE MANAGEMENT

2.6.1 GENERAL DESCRIPTION

In previous WP5 deliverable D5.2 [1], we examined a distributed radio-resource management scheme, which allowed subnetworks to autonomously adjust their transmit power using locally learned information and limited exchanges with neighbouring subnetworks. More specifically, a PoC was presented that brought the simulation-only work of WP4 deliverables [1] into an over-the-air testbed. In this setup, the 6G parent network first performed centralized offline training to capture typical interference patterns and produced a Graph-Neural-Network (GNN)

model for power control. The trained weights were then “transferred” to each subnetwork, so that, during operation, every node could run the model locally, using only its own sensing data, and exchange a few compact state vectors with its neighbours. This allowed the subnetworks to allocate power efficiently without continuous central coordination and adapt quickly when the topology changed, for example, when a sub-network joined or left. Both the Poc and the simulation work specifically targeted the "Subnetwork Coexistence in Factory Hall" use case defined in WP2 [1], which involves multiple industrial subnetworks operating within a shared physical space. In such environments, managing mutual interference is critical to ensure reliable communication - a challenge that becomes more pronounced in dense deployments. The main KPI used for evaluation is throughput, as defined in the description of this use case, where ensuring sufficient data rates under overlapping spectrum conditions is a primary performance requirement. Finally, the model-handoff mechanism appearing in this scheme, also shaped part of the architectural work in WP2 [1] , where AI/ML RRM model training and transfer -from the 6G parent to its subnetworks - was recognized as a core element of sub-network integration within the 6G ecosystem.

2.6.2 ENHANCED/NEW FUNCTIONALITY VS PREVIOUS VERSION

In D5.2 [1], the GNN-based interference management model was trained using the same parameter settings as those used in prior simulation-based evaluations from WP4. While this approach validated the concept in theory, it did not fully reflect the characteristics of the SDR-based PoC environment. This choice was obvious from the results, where it was observed that the GNN model performed well in low-to-medium SINR conditions -consistent with the simulation settings -but provided limited or no performance gains in high-SINR scenarios, where interference between subnetworks was minimal. As a result, the D5.3 implementation introduces revised training with parameters tailored to the measured operating conditions of the PoC platform.

Additionally, the updated PoC provides a more detailed breakdown of the power control scheme’s impact on both the individual link level and the network level. Unlike D5.2, where only average or harmonic mean results were highlighted, the new results illustrate how each subnetwork’s throughput is directly affected by the models’ decisions across different SINR combinations. The results clearly show how power is redistributed to favour the weaker link in asymmetric scenarios, confirming that the model improves network fairness.

2.6.3 HIGH LEVEL ARCHITECTURE CHANGES/UPDATES

As described in D5.2, we deployed the proposed RRM solution in a functional environment using the Open5GS and srsRAN software suites alongside USRP SDR hardware. Open5GS is an open-source core network implementation designed for constructing and managing NR/LTE mobile networks, [29]. Similarly, srsRAN is an open-source platform that allows researchers, developers, and telecom professionals to implement and experiment with LTE and 5G protocols using SDRs or a fully software-based setup, [30]. It offers a flexible framework for developing, testing, and deploying cellular network technologies, making it a valuable asset for wireless communication research. Its modular architecture and extensive documentation ensure accessibility for both academic and industry users. Ettus Research, a National Instruments (NI) brand, is a prominent provider of SDR hardware and solutions. Best known for its USRP product line, Ettus Research supports researchers, engineers, and developers in prototyping and deploying wireless communication systems, with applications ranging from spectrum analysis to signal processing and 5G research [31].

All the software components are executed by a single powerful computer. Apart from the gNBs and UEs, instances of the srsRAN software suite, and the Open5GS core network running in a Docker environment, the computer also handles all Message Passing Neural Network (MPNN) functions, implemented by PyTorch [32], including two vertices - one for each pair - and a centralized scheduler process. The scheduler communicates with all MPNN vertices and coordinates them to enable organized message-passing broadcasts. Coordination is crucial, as the vertices may operate at different execution speeds, but power decisions are only valid when all messages have been received and aggregated. Data exchange between processes is managed through inter-process communication using the ZeroMQ library, which supports multiple programming languages and provides a lightweight, flexible messaging layer that abstracts the complexities of socket programming [33].

A general architecture of implementation with the processes is displayed in Figure 36(a). For our experiments, we utilized four B210 SDR devices forming two subnetworks. Each subnetwork includes one gNB cell forming an HC node and one UE as an SNE. Figure 36(b) displays the devices placed on a testbed and their respective roles. All SDRs devices were connected to the same computer using USB3 cables. At this point it is worth noting that only the downlink communication is taking place wireless with omnidirectional antennas of 15cm, while the uplink is happening through SMA cables. This way, the presented results on the following section are only influenced by the wireless interference presented on the downlink.

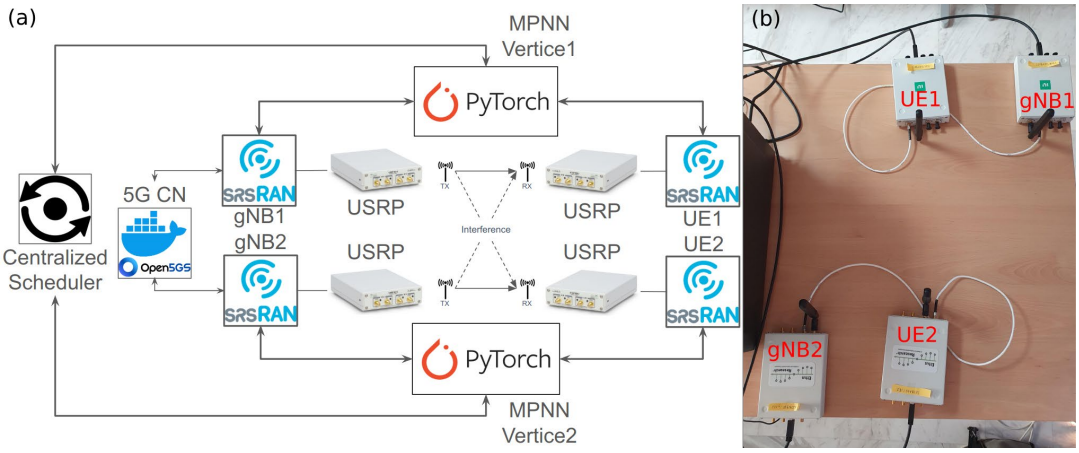

FIGURE 36. EXPERIMENTAL SETUP (A) AS AN ILLUSTRATION BETWEEN THE PROCESSES AND SDR DEVICES AND (B) IN AN ACTUAL ENVIRONMENT

2.6.4 FINAL DEMONSTRATOR SCENARIO DESCRIPTION UPDATES/CHANGES

For the demonstration of our proposed solution, we configured the SDR devices and the communication according to the parameters shown on the Table 7 . The operating carrier frequency was selected within the unlicensed frequency bands due to regulatory considerations, while the antenna gains were tuned to allow a wide range of SNRs as the SNEs can have maximum 1.5m distance from their serving HCs. The communication parameters were typical for mobile wireless communication.

TABLE 7. WIRELESS PARAMETERS

RF Parameter	**Value**
Carrier Band	N46 (5.4GHz)
Sample Rate	11.52 MS/S
Antenna Length	15cm
Tx/Rx RF ChainGain	58dB/40dB
Max SNE Distance	1.5m
Communication Parameter	**Value**
Bandwidth	10MHz
SCS	15KHz
Symbol/Slots	14/10
Frame Period	10ms
Cyclic Prefix	Normal

Regarding the neural network training, the process is the same as the one presented in the D4.3, [1]. Some hyper-parameters were altered as it was found that they provided better performance on the real demonstration scenario. The list of the final parameters is presented on the Table 8.

TABLE 8. NEURAL NETWORK'S HYPER-PARAMETERS.

Parameter	Value
Message MLP Φ	{9, 32, 32, 1}
State Update MLP U	{10, 32, 8}
Output MLP Ω	{8, 16, 1}
Embedding size	8
Epochs	100
Batch Size	10
Initial Learning Rate (LR)	0.002
LR decay factor	0.9
LR decay step	10
Optimizer	Adam

Just like the previous deliverable D5.2 [1], we utilized the iperf3 tool for traffic generation [34]. We knew that for the given communication parameters the theoretical maximum rate, when there is no interference on the communication channel, is around 45Mbps. Thus, the requested bandwidth on the iperf3 tool to generate traffic was 50Mbps. As expected, each pair was unable to achieve such a rate due to noise and interference with each other, although the srsRAN gNBs (HC nodes) were configured to perform Outer-Loop-Link-Adaptation (OLLA) algorithm. That said, each pair was constantly trying to achieve the maximum possible rate for given channel conditions, by changing the Modulation Coding Scheme (MCS) and retaining less than 10% BLER.

2.6.5 RESULTS

This section presents the measurement results of the proposed Proof of Concept (PoC) and compares them to a baseline bare-metal setup without power control for interference mitigation. Fifteen different test scenarios were evaluated, each reflecting distinct channel conditions between two subnetwork pairs. Desired channel conditions were achieved by monitoring the reported SNR from the UEs while slowly adjusting their positions in space. Increasing the SINR involved moving the UE closer to its serving gNB, while decreasing it meant placing it farther away and nearer to the interfering cell. The tested SNR values ranged from 3 dB to 15 dB.

The measurements yielded the following observations: generally, as each pair's SNR increased, so did its throughput, consistent with expectations. Higher-quality channels produced higher data rates. Notably, the greater the disparity in channel conditions between the two pairs, the larger the observed performance gains. Improvements ranged from 2% up to 42.2%, validating the implementation's effectiveness in dynamically adjusting power levels -reducing power for the stronger pair and enhancing conditions in favour to the weaker link.

Conversely, in scenarios where both pairs experienced similar channel conditions (e.g., combinations like 3-3, 6-6, 9-9, 12-12, and 15-15), no performance improvement was expected since the optimal strategy in such cases is Equal Power Allocation (EPA). Yet, the implementation underperformed in these instances, with the worst-case showing a 6.7% loss. This drop in performance likely stems from suboptimal training of the neural network. Figure 37 displays throughput results for each pair under both the GNN-based power control and EPA strategy. The intended behaviour is for the stronger pair under EPA to reduce its throughput under GNN by lowering transmission power, thus allowing the weaker pair to achieve better performance. This trade-off is evident in scenarios such as 3-9, 3-12, 3-15, and 9-15. Performance gains from GNN over EPA are further illustrated in Figure

38, where each matrix entry shows the percentage improvement in the harmonic meaning of both pairs' throughput. The harmonic mean is used to promote fairness, as it emphasizes lower throughput values more strongly.

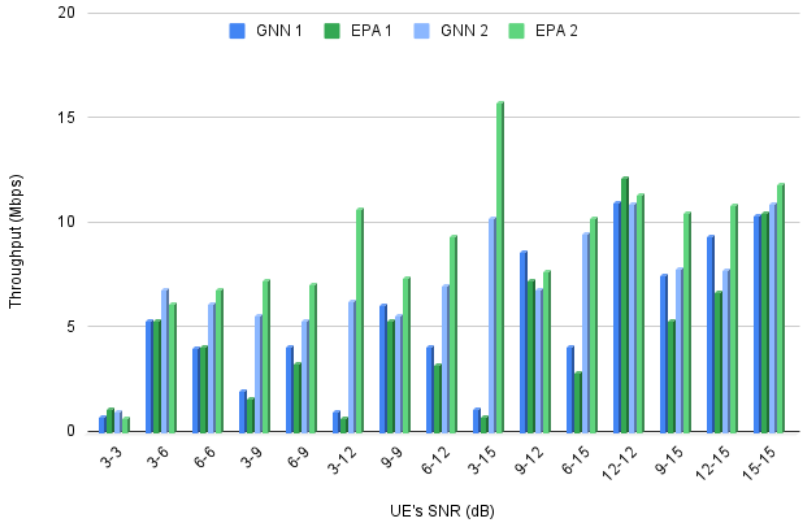

FIGURE 37. THROUGHPUT AS REPORTED BY EACH UE OVER DIFFERENT SNR COMBINATIONS FOR EACH UE WHEN APPLYING GNN AND EPA POWER CONTROL POLICY.

SNR2 / SNR1	3dB	6dB	9dB	12dB	15dB
3dB	-1.69%	4.52%	10.06%	32.37%	42.21%
6dB		-5.26%	3.21%	7.70%	29.39%
9dB			-5.50%	2.07%	8.56%
12dB				-6.73%	2.40%
15dB					-4.32%

FIGURE 38. IMPROVEMENT OF THE NETWORK'S HARMONIC MEAN THROUGHPUT FOR DIFFERENT CHANNEL CONDITION COMBINATION

2.6.6 ACHIEVED RESULTS AND WP2-3-4 THEORETICAL RESULTS

Simulation results of the proposed framework were presented in WP4 deliverables, where it was evaluated under controlled conditions using synthetic channel models. However, the outcomes of the simulations and the PoC are not directly comparable due to fundamental differences in both the radio environment and the way the message passing mechanism is implemented. In the simulations a generalized short-range channel model is used that included a wide variety of propagation conditions, such as multipath and non-line-of-sight components, to evaluate the algorithm's robustness across diverse deployment scenarios. In this setting, each subnetwork transmitted its local state using CSI-RS signals, which were aggregated over the simulated channel before being processed by the receiving node. Given this set-up, the results of WP4 demonstrated an average throughput gain of approximately 7%, with a maximum gain of 13% observed in scenarios with highly asymmetric channel

conditions. These evaluations confirmed the effectiveness of the GNN-based distributed power control scheme in improving throughput fairness, particularly in scenarios with asymmetric SINR conditions and multiple interfering subnetworks. While the PoC implementation in D5.3 follows the same core methodology, it operates under different assumptions and constraints, which lead to some divergence in observed outcomes. In the PoC, only two subnetworks were deployed due to hardware limitations, and communication took place under short-range, predominantly line-of-sight conditions with minimal external interference. The message passing mechanism, which is central to the distributed power control scheme, was embedded into the SSB bursts. The reason behind this choice lies in the practical requirements of the testbed regarding synchronization between the subnetworks and coordinated operation of the message passing procedure.

Furthermore, to reflect the specific conditions of the testbed, the GNN model used in the PoC was retrained with hyperparameters tuned to the measured channel characteristics. Under these conditions, the PoC demonstrated throughput improvements with an average gain of 17%. The observed gains ranged from 2% to a maximum of 42%, depending on the degree of asymmetry between the two subnetworks. The highest improvements were recorded when one subnetwork operated under significantly lower SINR than the other, allowing the power control algorithm to shift resources toward the weaker link. This mirrors the same behavioural trend observed in simulation: the algorithm improves throughput fairness under asymmetric channel conditions, while offering little benefit when links are already balanced. The higher gains observed in the PoC are explained by the specific characteristics of the testbed environment, such as short-range line-of-sight channels and the retraining of the model using measured data. While the test conditions differ from those used in simulation, the consistency in performance trends confirms the practical viability of the proposed approach. These results demonstrate that the distributed power control scheme remains effective under real-world constraints and provide a foundation for future extensions toward more complex subnetwork deployment scenarios.

2.6.7 CONCLUSIONS

The presented final results from the updated PoC confirm and extend the initial findings presented in deliverable D5.2. The implementation demonstrated significant throughput improvements in a wide range of asymmetric interference scenarios, with observed gains reaching up to 42% over the baseline. In scenarios with more balanced SINR conditions, where the optimal policy resembles equal power allocation, the benefits of GNN-based control were limited or, in some cases, negative -highlighting areas for further investigation. Nevertheless, the overall results validate the potential of the proposed scheme for improving spectral efficiency in realistic subnetwork coexistence settings and provide clearer insight into its behaviour across different interference regimes.

During the PoC development, several practical challenges were identified. The most critical one was the absence of synchronization between HC nodes, which led to the use of SSB-based message passing instead of the simultaneous CSI-RS-based over-the-air aggregation assumed in WP4. Furthermore, the limited hardware setup constrained the evaluation to two subnetworks, restricting insights into scalability. Future steps will focus on expanding the number of subnetworks and enabling true over-the-air message aggregation, allowing for a more comprehensive assessment of the solution under broader and more complex interference scenarios.

2.7 JAMMING DETECTION USING ANOMALY DETECTION METHODS

2.7.1 GENERAL DESCRIPTION

In Deliverable D2.4 of WP2 [1], two key architectural components were introduced that enable sensing operations in 6G, and to support the integration of subnetworks to the parent 6G network. These two components are the Sensing Coordination Function (SCF) and Sensing Analytics Function (SAF). The SCF coordinates the sensing activities, e.g., when to start or stop sensing or which sensing entities (UEs or BSs) are involved in sensing activities. The SAF collects sensing data from sensing entities and comprises an analytics engine capable of complex sensing data processing and/or fusion, and conversion of sensing data into sensing results. Sensing data is broadly speaking the low-level data generated from the radio interface out of which sensing results can be generated. For a full

description of sensing results, consult D5.2 [1]. For the purposes of this PoC, positioning of a target object is the considered sensing result. This is important in many cases, but the *Virtual ECU: In-vehicle sensor data and functions processing at the 6G network edge*, is chosen for this work, due to its focus on the collaboration between the parent 6G and the subnetworks at compute level.

In this deliverable, the focus is on the integration of the in-vehicle subnetwork with the parent 6G network, where processing of sensing data can take place in both networks. A range detection algorithm was developed for extraction of relative distances from sensing data, and that algorithm runs in an Edge cloud co-located with the 6G CN, more specifically at the SAF, in line with D2.4 [1]. The algorithm operating in the SAF will have its performance impaired in the presence of malicious or legitimate jammers, causing the parent 6G network to have to take action and reject the transmission of the sensing data from the subnetwork due to various reasons, such as poor quality or no timely delivery of sensing data not allowing the SAF to determine sensing results that meet the sensing service KPIs (positing accuracy), or suspected replaying of existing sensing data caused by malicious jammers. Following that action, the subnetwork will be required to use its own compute resources. There are several possibilities for the algorithm design presented in D4.3 [1], and the enablers for the collaborative compute functions herein are described in D4.4 [1].

2.7.2 ENHANCED/NEW FUNCTIONALITY VS PREVIOUS VERSION

This version of the PoC is enhanced in relation to D5.2 [1] in the following aspects: USRPs have been programmed for the transmission of sensing pilot signals, in a single transmitter and single receiver setup. The receiver has been programmed to generate files that encode CIR as measured by the sensing receiver. The receiver has then been programmed to stream these files to the SAF. The distance or range detection algorithm has been developed, and can be executed both at the SAF, or at the vehicle subnetwork. Finally, the protocols described in D4.4 [1] were implemented, enabling distributed compute support between subnetwork and parent 6G network.

2.7.3 HIGH LEVEL ARCHITECTURE CHANGES/UPDATES

The high-level functional blocks and logical flows for the procedures described above are depicted in Figure 39 and Figure 40. The sensing transmitter generates a sensing RF pilot signal. This is done resorting to the GNU radio SW. Then, because the setup is cabled, a channel model was used to emulate the channel. The sensing RF pilot then flows through a channel model using a workstation with GNU radio installed. The sensing pilot is then passed on to the USRPs, that transmit it over to the sensing receiver, via a cabled setup. The sensing data is transported using a dedicated format. The sensing receiver has its own USRPs, programmed to receive the files. The data is then processed to extract meaningful sensing results.

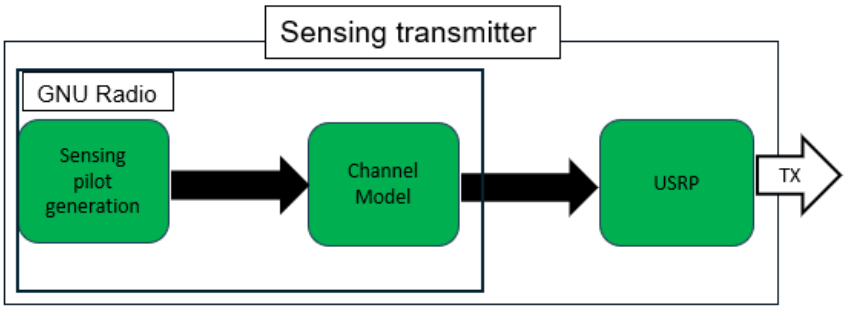

FIGURE 39. LOGICAL FLOWS FOR THE SENSING TRANSMITTER

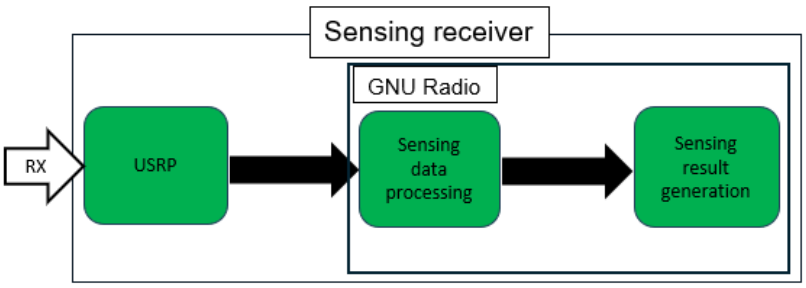

FIGURE 40. LOGICAL FLOWS FOR THE SENSING RECEIVER

Table 9 and Table 10 contain the detailed list of all hardware and software used in this PoC, respectively.

TABLE 9. HARDWARE LIST USED IN THE POC

Name	Info
Dell Precision 7820 Tower	Intel Xeon Gold 5122, 64 GB memory, Intel X710 10 GbE NIC
Dell Precision 7810 Tower	Intel Xeon Silver 4114, 64 GB memory, Intel X710 10 GbE NIC
USRP N310 (2)	ZYNQ-7100, 4 CHANNELS, 10 MHZ - 6 GHZ, 10 GbE NIC

TABLE 10. SOFTWARE LIST USED IN THE POC

Name	Info
Ubuntu 22.04	OS
GNU Radio 3.10.7.0	SDR programming and signal processing
UHD 4.7.0.0	A user-space library that communicates with and controls USRPs

2.7.4 FINAL DEMONSTRATOR SCENARIO DESCRIPTION UPDATES/CHANGES

The in-vehicle subnetwork is authenticated and authorized for sensing service within the 6G parent network. The sensing subnetwork is performing bi-static sensing, i.e., the sensing receiver picks up the sensing pilots emitted by the transmitter. The receiver produces sensing data files based on the CIR in a continuous fashion, and streams these to the SAF. The SAF runs the range detection algorithm on the received sensing data. This is considered to be a normal functioning operation of the PoC. Next, an anomaly is simulated by injecting artificial noise into the sensing signal. This will not influence the processing at the sensing receiver, that will continue to stream this data to the SAF. The SAF however, will determine that there is a problem with the sensing data, and that the noise levels make it unfeasible to produce accurate sensing results, as previously described in D5.2 [1].

What happens next is a transfer of functionality from the SAF to the subnetwork, whereby the subnetwork is required to perform processing over sensing data, while the interference conditions are still observable. This is done resorting to similar signalling presented in D4.4 [1]. However, the SAF is considered to have unlimited compute resources, but the vehicle subnetwork does not – the processing resources are much lower there, which can lead to delays in the computation of other in-vehicle processing tasks, that are totally unrelated to the sensing operation. For that reason, we present a study on compute capabilities and processing time for the sensing data being generated. This may pose challenges to the vehicle, but it is the most straightforward way to deal with an interference affected channel. The interruption of the streaming of sensing data to the SAF is therefore, a part of the interference mitigation process. The underlying argument of processing sensing data within a subnet before sharing the outcome (i.e. sensing results) with the parent network yields advantages, i.e., QoS requirements for

sharing sensing data is higher in throughput and more stringent in terms of latency and jitter due to the sensitivity of the SAF's algorithm for fusion and processing purposes.

2.7.5 RESULTS

2.7.5.1 BASELINE RESULTS UPDATE

Results are presented sequentially in this section. First, Figure 41 shows an example of distance/range as a sensing result being extracted from sensing data. If there are objects around between the sensing transmitter and receiver, the sensing signals will be reflected by them and will be detected by the sensing receiver. In Figure 41, the signal is propagated without the presence of any jamming, and the peaks detected from reflections are easily visible and marked with crosses. In Figure 42, peak detection does not work because of the noise introduced in the channel.

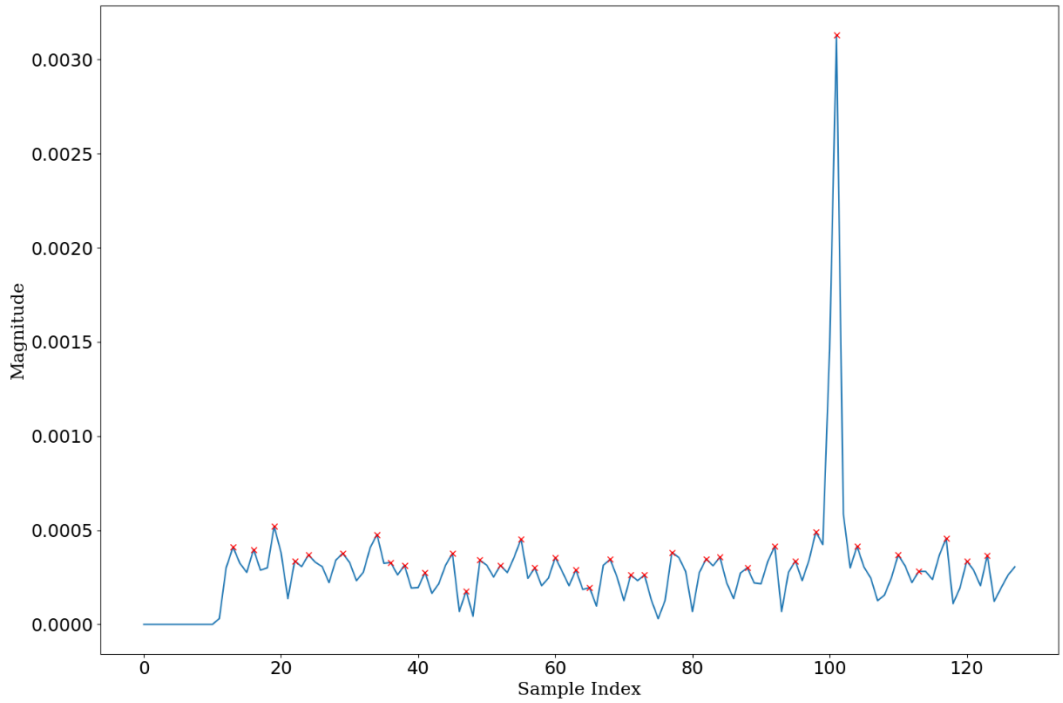

FIGURE 41. RANGE DETECTION ALGORITHM IN SAF FROM SENSING DATA - NO SIGNIFICANT INTERFERENCE

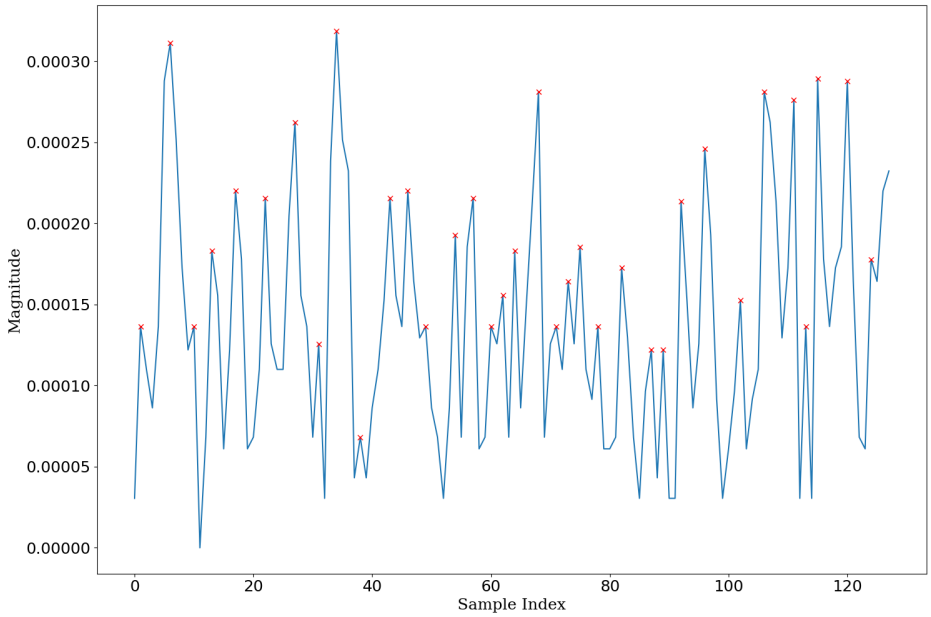

FIGURE 42. RANGE DETECTION ALGORITHM IN SAF FROM SENSING DATA - WITH INJECTED NOISE

Once this happens, a subnetwork switch over is performed. The vehicle subnetwork is configured to interrupt sensing data transmission and is required to process sensing data for the duration of the interference. Hence, it is important to understand the load that this processing activity will add to the in-vehicle subnetwork.

The sensing data is generated at the sensing receiver and the SCF can instruct the sensing receiver to change the sampling rate for purposes such as energy consumption, a change in accuracy requirements or compute availability constraints. The sampling rate is the periodicity at which the sensing data is being determined from the measurements. This is an effective mechanism that reduces the computational burden, and that is consistent with the presence of the interferer, where the sensing result will never be accurate anyway. Hence, reducing this sampling rate and compute sensing results from sensing data is quite suitable for understanding if the interference is still present.

Figure 43 shows the average computational processing time using a box and whisker plot including the lower and upper quartile, extreme quartile and outliers. The computational processing time is collected at the sensing receiver when determining the sensing data for different sampling rates.

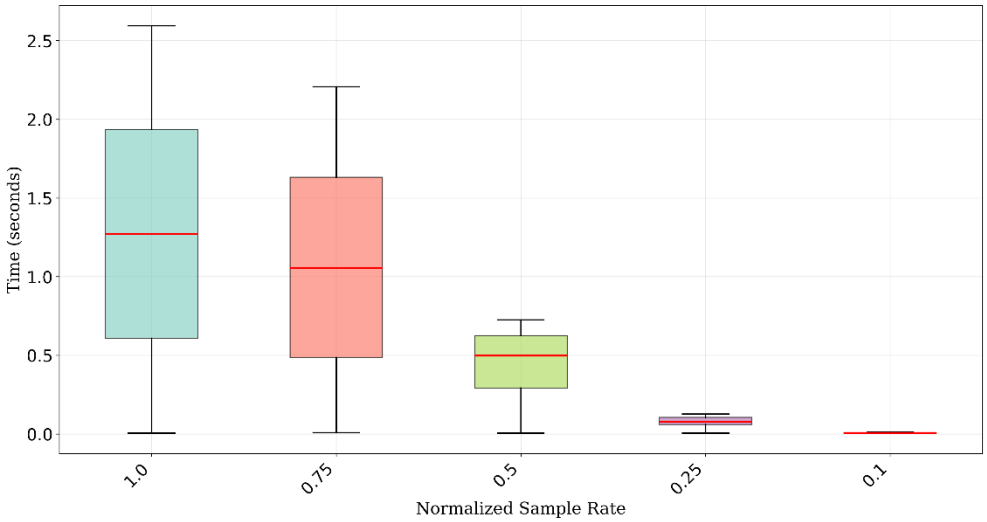

FIGURE 43. AVERAGE COMPUTATIONAL PROCESSING TIME FOR DIFFERENT SAMPLING RATES TO PRODUCE SENSING DATA

Results show prohibitive compute time if the same sampling rate used by the SAF is also used in the subnetwork. An expected descending compute time trend is observed as the sampling rate decreases. This is expected because there is less data to process. As a result, different sampling rates settings can be configured by the SCF, as indicated in the general description of this PoC.

Finally, the same CPU with different processing profiles is tested against different sampling rates and results are shown in Figure 44. The profiles used are on demand and performance, and are used to simulate different subnetwork nodes, e.g., and HC and an LC. The x and y axis are normalized to the highest average compute time recorded, and the highest sample rate, respectively.

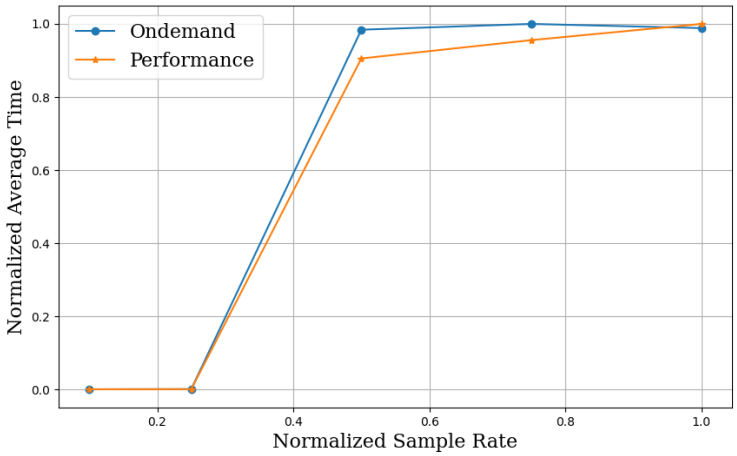

FIGURE 44. COMPARISON BETWEEN TWO CPU PROFILES, FOR DIFFERENT SAMPLING RATES FOR THE SENSING DATA

Results show a difference of roughly 10% more compute time for a CPU optimized for computing tasks on demand based on input buffer, against a CPU tuned for performance, which is always readily available for computing at maximum performance. This was used for simulating two CPUs with different compute capabilities, similar to an HC and LC node. Results are normalized but together with the ones in Figure 43 can be used to estimate the absolute value for computing time.

The sampling rate can be seen as an adjustable parameter that can be tailored to different subnetworks, due to its overall different compute capabilities. In that case, a value that is low enough and that suits the subnetwork nodes

can be chosen and fixed. The subnetwork compute ratio may also change over time. Hence, the sampling rate can be negotiated between SAF and subnetwork, and this can happen once or in run time, while the subnetwork processes the data. Another approach is for the subnetwork to require to process at a certain sampling rate, in order to be able to accept the SAF's request. Either way, the key point is that it becomes an important tunable parameter and can be included as part of the procedures for the distributed compute fabric presented in D4.4 [1].

No formal analysis is made comparing different anomaly detection algorithms presented in D4.3 [1], because the usage of a different algorithm would result in similar results as the ones presented in this previous section, simply scaling differently.

2.7.5.2 CONCLUSION

This PoC demonstrator enables the collection of FR-based sensing data and the production of sensing results such as range/distance, that are important to many use cases and particularly useful for mobile subnetworks, such as vehicles. A range/distance detection algorithm is implemented that can be executed at both SAF in the parent network and the SAF in the subnetwork. The results show that the compute capabilities for sensing data generation inside the sensing receiver and processing of sensing data into results by the SAF are vital in supporting constrained subnetwork nodes. Additionally, continuous generation of sensing data at the highest possible sampling rate at the sensing receiver and streaming of all sensing data to the SAF for processing into sensing results increases the chances of sampling underuns (in the sensing receiver) and processing overuns (in the SAF). The PoC addresses this challenge by providing findings that demonstrate that the availability of compute resources has a direct impact on the ability to sample the radio interface to produce sensing data accurately (i.e. timely). Moreover, the PoC has shown that range detection algorithms significantly suffer from noisy environments increasing the likelihood for failed range detection. Thus, the insights from the PoC forms input to future work towards efficient solutions that are suitable for compute-constrained subnetworks, implementing functionalities that are presented in D4.3 and D4.4 [1], and establishes a clear line of coordination options of the SCF which is located in the parent 6G network.

3 CONNECTION WITH WP2 USE-CASES

After summarizing the presentation of all PoCs and their achieved results, we would like to present an overview of the achievements related to the use-cases defined in WP2, in order to clearly present what the real-world benefits from the presented TCs can be in relevant use-cases identified in WP2.

In Table 11, we summarize the implemented PoCs and the related WP2 use-cases they are directly targeting as well as the relevant KPIs defined as crucial in these use-cases. We then choose the 2 most important KPIs and present the overall performance enhancements achieved compared to the SoTA baseline.

TABLE 11. SUMMARIZE TABLE OF POC PERFORMANCE GAINS VS WP2 USE-CASES REQUIREMENTS

POCs	**WP2 USECASE**	**Related KPIs**			
		KPI	**% gain**	**KPI**	**% gain**
Low latency channel emulator	Channel modelling	End-to-end latency	2× reduction vs initial value	Performance improvement in critical DSP processing blocks	3×
Jammer Resilient PHY	Robot control	Throughput	Less than 1% loss vs non jammed system	N/A	NA
Latency aware MAC	Industrial use cases Robot control	Latency	Up to 98.6% lower than SoTA technology baseline	Jitter	Up to 99.6% lower
Intra-subnetwork macro-diversity	Industrial use cases Robot control	Reliability	PLR of $2.0 \cdot 10^{-9}$ is up to 5000 times lower than 5G URLLC	Cycle time	10x faster compared to RTT in 5G
Wireless in-vehicle E/E architecture	Inter-subnetwork Coordination: Collaboration/interference/RRM between subnetworks in intra/inter-vehicle communications	Latency	Similar performance compared to automotive ethernet and CAN	Throughput	Similar performance compared to automotive ethernet and CAN
Centr/Distr interference management	Subnetworks Swarms: Subnetwork Coexistence in Factory Hall	Throughput	17% on average and up to 42% gain compared to EPA baseline	N/A	N/A
Jamming Detection using anomaly detection	Virtual ECU: In-vehicle sensor data and functions processing at the 6G network edge	Determinism, reliability	Compute time as a direct function of the sampling rate with a 0.0s average and percentiles compute time for a normalized sampling rate of 0.1	-	-

As can be seen in all selected PoCs there are significant gains or achieved functionalities that are beyond the state-of-the-art of wireless communications. Being able to provide gains from 1.5X, 2X all the way to 98.6% better than the baseline in some cases is a clear sign that the tested TCs in the selected PoC of WP5 have the potential to be real enablers and even game-changers in the effort to move to the era of in-X subnetworks. Subnetworks offering similar performance to wired ethernet CAN in the case of collaborative RRM within vehicles PoC, as well as diminish throughput losses down to 1% when subnetworks are interfered from external systems, mitigating the interference impact in the running subnetwork are also quite important achievements that were validated in real life through the implementation of our PoCs.

Overall, we can conclude that the selected PoCs were able to demonstrate their potential in real life experiments, with environment and scenario designs aiming to be similar to identified use-cases in WP2, allowing us to have a clear view on the applicability of these TCs in relevant use-cases and their expected real-life impact in performance.

4 CONCLUSIONS

In this deliverable, we have presented the final version and results of the 6G-SHINE PoCs that are defined as the target TCs to be implemented in WP5. Specifically, we targeted seven PoCs: Low latency channel emulator (addressing TC2), Latency aware MAC access (addressing TC11), Jamming resilient PHY (addressing TC6), Intra-subnetwork macro-diversity (addressing TC8), Wireless in-vehicle E/E Architecture (addressing TC13), Centralized/distributed interference management (address both TC12 and TC13), Jamming detection using anomaly detection (addressing TC15). These PoCs successfully demonstrate that the selected technology components have reached Technology Readiness Level 4.

For each PoC, we have presented the demonstrator scenario, a baseline which the PoC outperformed, a description of the PoC including its enhanced functionality and the use cases defined in WP2 where the TC is an important enabler for. We further compared the performance achieved in the PoCs to the previous PoC version if possible as well as to the theoretical/simulation results from WP2/3/4 where the relevant TCs were proposed and analysed in detail. This allowed the validation of results presented in these WPs against what reality says in WP5. Performance was verified against the identified baseline KPIs and against the requirements of related use-cases. The results presented clearly demonstrate that the selected TCs implemented as PoC in WP5 offer significant performance benefits or novel functionalities that go beyond the current state of the art. These outcomes are consistent with the theoretical advantages anticipated from earlier simulation and analytical evaluations carried out in WPs 2, 3, and 4. Together, they reinforce the conclusion that many of the proposed TCs have the potential to be true game changers in the design and deployment of in-X subnetworks, paving the way for the realisation of the envisioned use cases.

Finally, possible shortcomings in available HW/SW or other implementation issues were identified that need to be addressed in the near future in order to make the mass adoption and production of relevant TCs in the in-X subnetworks of tomorrow.

5 REFERENCES

[1] 6G-SHINE project deliverables: https://6gshine.eu/deliverables-ii/

[2] "F8800B PROPSIM F64 Radio Channel Emulator," Keysight technologies, [Online]. Available: https://www.keysight.com/us/en/assets/3122-2164/data-sheets/F8800B-PROPSIM-F64-Radio-Channel-Emulator.pdf.

[3] Degli-Esposti, V.; Vitucci, E. M.; Di Renzo, M.; Tretyakov, S. A., "Reradiation and Scattering From a Reconfigurable Intelligent Surface: A General Macroscopic Model," IEEE Transactions on Antennas and Propagation, vol. 70, no. 10, pp. 8691-8706, October 2022.

[4] "F9860000A Channel Studio GCM," Keysight Technologies, [Online]. Available: https://www.keysight.com/us/en/product/F9860000A/channel-studio.htmlhttps://www.keysight.com/us/en/product/F9860000A/channel-studio.html.

[5] "E7515B UXM 5G Wireless Test Platform," Keysight Technologies, [Online]. Available: https://www.keysight.com/us/en/product/E7515B/uxm-5g-wireless-test-platform.html.

[6] "S8708A 5G Advanced Performance Test Toolset," Keysight Technologies, [Online]. Available: https://www.keysight.com/us/en/product/S8708A/s8708a-advanced-performance-test-toolset.html.

[7] IEEE Standard for Information Technology--Telecommunications and Information Exchange between Systems Local and Metropolitan Area Networks--Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 1: Enhancements for High-Efficiency WLAN (2021) https://standards.ieee.org/ieee/802.11ax/7180/

[8] European Telecommunications Standards Institute (ETSI): ETSI TR 103 665.Technical Report TR 103 665 V1.1.1 (May 2021). https://www.etsi.org/deliver/etsitr/103600103699/103665/01.01.0160/tr103665v010101p.pdf

[9] Federal Communications Commission (FCC): FCC Title 47 CFR Part 15.407: Subpart E - Unlicensed National Information Infrastructure (U-NII) Devices (2023). https://www.ecfr.gov/current/title-47/chapter-I/subchapter-A/part-15/subpart-E/section-15.407

[10] Thijs Havinga and Xianjun Jiao and Wei Liu and Baiheng Chen and Adnan Shahid and Ingrid Moerman: Wi-Fi 6 Cross-Technology Interference Detection and Mitigation by OFDMA: an Experimental Study. (2025) https://arxiv.org/abs/2503.05429

[11] Yao, Junmei and Lou, Wei and Xie, Ruitao and Jiao, Xianlong and Wu, Kaishun: Mitigating Cross-Technology Interference Through Fast Signal Identification. In: IEEE Transactions on Vehicular Technology 72(2), pp2521-2534 (2023). https://10.1109/TVT.2022.3213663

[12] De Valck, Peter and Tytgat, Lieven and Moerman, Ingrid and Demeester, Piet: Coexistence Aware Clear Channel Assessment. In: Wireless Sensor Networks, pp. 165–178. Springer, Berlin, Heidelberg (2013)

[13] Jiao, Xianjun and Liu, Wei and Mehari, Michael and Aslam, Muhammad and Moerman, Ingrid: openwifi: a free and open-source IEEE802. 11 SDR implementation on SoC . In: 2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring), pp1-2, 2020.

[14] Muhammad Aslam and Xianjun Jiao and Wei Liu and Michael Mehari and Thijs Havinga and Ingrid Moerman: A novel hardware efficient design for IEEE 802.11ax compliant OFDMA transceiver. In: Computer Communications, v219, pp173-181, https://doi.org/10.1016/j.comcom.2024.03.006

[15] Havinga, Thijs and Jiao, Xianjun and Liu, Wei and Moerman, Ingrid: Experimental Study Towards Efficient Interference Avoidance Using Wi-Fi 6 OFDMA on SDR. In : EEE INFOCOM 2024 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), 2024, doi: 10.1109/INFOCOMWKSHPS61880.2024.10620761

[16] Tjijs Havinga, Xianjun Jiao, Wei Liu Baiheng Chen, Robbe Gaeremynck, Ingrid Moerman: Cross-Technology Interference Awareness for Multi-User OFDMA Scheduling in IEEE 802.11ax, submitted to Ad Hoc Networks, preprint via http://ssrn.com/abstract=5266345.

[17] https://www.xilinx.com/products/boards-and-kits/ek-u1-zcu102-g.html

[18] https://www.analog.com/en/resources/evaluation-hardware-and-software/evaluation-boards-kits/eval-ad-fmcomms2.html

[19] https://www.intel.com/content/www/us/en/products/sku/189347/intel-Wi-Fi-6-ax200-gig/specifications.html

[20] ETSI Standard EN 300 328 V2.2.2, HTTPS://WWW.ETSI.ORG/DELIVER/ETSI_EN/300300_300399/300328/02.02.02_60/EN_300328V020202P.PDF
[21] OpenWRT. "Welcome to the OpenWRT Project." Accessed October 28, 2024. https://openwrt.org/
[22] D. Bankov, E. Khorov, A. Lyakhov and M. Sandal, "Enabling Low Latency Communications in Wi-Fi Networks," 2018 IEEE 29th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC).
[23] China Automotive Vision Industry Report, 2022. Accessed June 13, 2025. http://www.researchinchina.com/Htmls/Report/2022/72819.html
[24] International SparkLink Alliance. Accessed June 13, 2025. https://sparklink.org.cn/en/
[25] V. -P. Bui, D. Abode, P. M. de Sant Ana, K. Muthineni, S. R. Pandey and P. Popovski, "Digital Twin of Industrial Networked Control System based on Value of Information," GLOBECOM 2024 - 2024 IEEE Global Communications Conference, 2024.
[26] Bosch Mobility. Nachbereichkameras. Accessed June 13. https://www.bosch-mobility.com/de/loesungen/kamera/nahbereichskamera/
[27] D. Abode, P. M. d. S. Ana, R. Adeogun, A. Artemenko and G. Berardinelli, "Goal-Oriented Interference Coordination in 6G In-Factory Subnetworks," in IEEE Journal on Selected Areas in Communications, 2025.
[28] C. Kinzig, I. Cortés, C. Fernández and M. Lauer, "Real-time Seamless Image Stitching in Autonomous Driving," *2022 25th International Conference on Information Fusion (FUSION)*, Linköping, Sweden, 2022.
[29] Open5GS. Available at: https://open5gs.org/ (Accessed: 22 May 2025).
[30] srsRAN project. Available at: https://www.srsran.com/ (Accessed: 22 May 2025).
[31] USRP B210 USB Software Defined Radio (SDR), Ettus Research. Available at: https://www.ettus.com/all#products/ub210-kit/ (Accessed: 22 May 2025).
[32] PyTorch. Available at: https://pytorch.org/ (Accessed: 22 May 2025).
[33] ZeroMQ. Available at: https://zeromq.org/ (Accessed: 22 May 2025).
[34] iPerf3. Available at: https://iperf.fr/ (Accessed: 22 May 2025).